

Discussion Paper Series

IZA DP No. 18778

July 2026

Labor Market Tightness and Hiring Outcomes: Evidence from Job Application Data

Lennart Ziegler

Central European University
and IZA@LISER

The IZA Discussion Paper Series (ISSN: 2365-9793) ("Series") is the primary platform for disseminating research produced within the framework of the IZA@LISER Network, an unincorporated international network of labour economists coordinated by the Luxembourg Institute of Socio-Economic Research (LISER). The Series is operated by LISER, a Luxembourg public establishment (établissement public) registered with the Luxembourg Business Registers under number J57, with its registered office at 11, Porte des Sciences, 4366 Esch-sur-Alzette, Grand Duchy of Luxembourg.

Any opinions expressed in this Series are solely those of the author(s). LISER accepts no responsibility or liability for the content of the contributions published herein. LISER adheres to the European Code of Conduct for Research Integrity. Contributions published in this Series present preliminary work intended to foster academic debate. They may be revised, are not definitive, and should be cited accordingly. Copyright remains with the author(s) unless otherwise indicated.



Labor Market Tightness and Hiring Outcomes: Evidence from Job Application Data*

Abstract

This study examines how the supply of job applicants shapes hiring outcomes of firms. Using Austrian matched vacancy-worker data, my analysis exploits variation in the number of job applications, inferred from referrals made by the public employment service. To address endogeneity concerns, I instrument the number of applications with the inflow of jobseekers from mass layoffs. A larger supply of applicants shortens vacancy filling duration and raises the job tenure of hires, consistent with improved worker-firm match quality. At the same time, more applications lower starting wages, which indicates that the employer-bargaining channel dominates the match-quality channel. A larger applicant pool also reinforces worker segregation: firms with high shares of male (native) workers hire even fewer female (foreign) workers when applications are abundant. This reflects employer selection rather than a shift in the composition of the applicant pool, which responds little to the supply shock.

JEL classification

J23, J62, J63, J71

Keywords

labor market tightness, vacancies, recruitment, worker-firm match quality, wage setting, worker segregation

Corresponding author

Lennart Ziegler

zieglerl@ceu.edu

* I thank David Card, David Green, Nikolaj Harmon, Andrea Ichino, Jeremias Klauui, Peter Kuhn, Giovanni Mastrobuoni, Michael Siegenthaler, Eduard Storm and Andrea Weber, as well as workshop and conference participants in Belgrade (Applied Economics Conference 2024), Toronto (EALE-SOLE Conference 2025), Marseille (LAGV Conference 2025), Zurich (KOF-IZA Workshop 2025), Nuremberg (IAB Workshop: Institutions and Frictions in the Labor Market 2025), Santiago de Chile (LM2C2 Workshop 2025), Edinburgh (Workshop on Employment Dynamics 2025) and Turin (LABORatorio R. Revelli Workshop 2026) for helpful comments and feedback.

1 Introduction

When a firm posts a vacancy, the number of suitable workers who apply can differ widely across otherwise similar jobs. This supply of applicants shapes how the firm hires. When many apply, the employer can fill the position quickly, wait for a worker who fits well and choose freely among candidates. When few apply, it must search longer, accept a weaker match or hire whoever is available. How readily firms find workers is a growing concern in the labor market. Many countries are expected to face stronger labor supply shortages in the upcoming decades as the retirement of the baby boomer generation reduces the size of the workforce. As not all sectors of the economy are equally affected by such shortages, worker movements across sectors and occupations will play a central role in the necessary adjustment of the labor market, requiring higher flexibility and more retraining of workers.

Stronger mobility can also lead to more mismatches between workers and firms. Despite retraining efforts, some hired workers might still not meet the necessary skill requirements for a given job. Such mismatches are costly. Workers whose education does not match the job's requirements tend to earn substantially less (Nordin et al., 2010; Altonji et al., 2016), and at the macroeconomic level occupational and sectoral mismatch is associated with productivity losses (Patterson et al., 2016) and higher unemployment (Şahin et al., 2014). This has severe implications for both workers and firms. Firms may become less productive and could be forced to lay off some workers again. Mismatched workers who are eventually laid off may struggle to reintegrate into the labor market. In contrast, higher mobility could benefit those workers who are qualified but face entry barriers due to labor market discrimination. If, for instance, the gender of applicants plays a major role in the firm's hiring decisions (Kuhn and Shen, 2023; Card et al., 2024), women find it more difficult to be hired for jobs typically held by men, and vice versa. Higher labor market tightness, i.e. a lower number of suitable job applicants per vacancy, might weaken the impact of such preferences in the hiring process. Similarly, workers of foreign nationality could benefit from worker shortages if employers discriminate against foreigners in loose labor markets.

This paper studies how the supply of applicants to a vacancy shapes these hiring outcomes. My analysis combines administrative data on job postings and job applica-

tions from the Austrian public employment service (AMS) with worker spells from the Austrian social security administration. To assist in the job search process, the employment service refers vacancies to registered jobseekers. My analysis uses the resulting applications as a proxy for the supply of applicants to each vacancy. I link vacancies and workers at the firm level, which is feasible for about half of all filled vacancies posted on the AMS job board.

Because the number of applicants may be related to unobserved features of the job and the firm, I instrument it with variation caused by mass layoffs. Specifically, the instrument is the inflow of registered jobseekers from mass layoffs in the same sector and district in the months before a vacancy is posted. This unexpected, sector- and region-specific inflow of jobseekers creates a positive supply shock for recruiting firms. Several robustness checks support the credibility of the layoff instrument. It is balanced on predetermined firm and vacancy characteristics, has no effect in placebo windows after vacancies are filled and is robust to alternative layoff thresholds. By linking the hired workers to their subsequent employment spells, I analyze how the supply of applicants affects vacancy filling, the match quality of hires, their pay and the gender and nationality composition of recruitment.

My estimation results show that a higher number of suitable applicants leads to shorter vacancy filling durations and longer job tenure of hires, consistent with an increase in worker-firm match quality. In contrast, more applications reduce starting wages. The negative effect is concentrated among higher-paid workers, whose pay is set with more discretion, while collective bargaining floors bind the pay of lower-paid workers in Austria. This suggests that a larger applicant pool strengthens the employer's position in wage setting, and that this bargaining effect dominates any wage gains from higher match quality. Stronger labor market tightness also attenuates firm segregation by gender and nationality: when the supply of suitable workers decreases, firms with higher shares of male workers hire relatively more women, and firms with higher shares of native workers hire relatively more foreigners. I show that this pattern reflects employer selection rather than the composition of the applicant pool, which does not change much when applications rise.

The role of labor market tightness for firms and workers has been a longstanding

question in labor economics, but little empirical evidence exists at the vacancy level. Macroeconomic research studying job creation and wage setting usually exploits time series data of aggregated labor market tightness at the national or state level (Shimer, 2005; Domash and Summers, 2022). Recently, several papers have started to explore variation at the firm and worker level. Davis et al. (2013) analyze hiring behavior at the establishment level using US data from the Job Openings and Labor Turnover Survey (JOLTS). Abraham et al. (2020) use micro data from the Current Population Survey to track worker flows across labor market states. Another viable data source in this context is online job postings, used by Azar et al. (2020) and Turrell et al. (2021) to study labor market tightness at the occupation level. Hoeck (2023) draws on linked firm-vacancy-worker data from Denmark to define firm-specific measures based on the occupational composition within firms. Klæui et al. (2024) use detailed click data from jobseekers and recruiters and find that, when the relative supply of suitable applicants is scarce, jobseekers are less likely and recruiters are more likely to search more broadly. By combining vacancy and job application data with employment spells, I extend the analysis of labor market tightness to examine the impact of vacancy-specific scarcity on realized hiring outcomes, including vacancy duration, wages and firm segregation.

This paper also connects to the literature in personnel economics studying the determinants and consequences of firms' hiring decisions (Oyer and Schaefer, 2011; Bloom and Van Reenen, 2011). A central aspect is the relationship between labor market tightness and worker compensation. When many firms compete for few workers, employers might increase wage offers to attract applicants, as in standard search-and-matching models (Pissarides, 2000) where jobseekers improve their bargaining position in tighter markets. Using the Austrian vacancy database, Mueller et al. (2024) show that vacancies fill faster when starting wages are higher, but the estimated elasticities are small. Cullen et al. (2024) analyze wage-posting data from a staffing platform and find that few firms adjust wage offers when faced with labor shortages. Pay-equity concerns can further limit such adjustments, as firms are reluctant to pay newly recruited workers more than their existing employees (Cullen and Perez-Truglia, 2022). My results instead indicate relatively strong employer bargaining power: a larger applicant pool lowers the wage of the hired worker despite an improved worker-firm match. This is consistent

with monopsonistic wage setting, in which greater applicant availability strengthens the employer's outside option (Manning, 2003; Azar et al., 2020).

A central insight of Becker (1957) is that increased competition makes it harder for firms to sustain discriminatory hiring behavior. Weber and Zulehner (2014) show for Austria that firms which discriminate against women are less likely to survive, especially in competitive industries. Exploiting business cycle variation, Biddle and Hamermesh (2013) show that tighter labor markets are associated with smaller gender and racial wage gaps. During the hiring stage, race-based discrimination against applicants is less prevalent when labor markets are tight (Baert et al., 2015; Doleac and Hansen, 2020). In line with the existing evidence on discriminatory hiring behavior, I find that firm segregation by gender and nationality decreases when vacancies receive fewer applications. Moreover, I provide evidence that this is driven by employer preferences rather than changes in the applicant pool.

2 Theoretical predictions

When labor markets are tight, employers find it harder to fill their vacancies: they can postpone hiring until a well-suited applicant arrives, or lower their requirements and accept a weaker match. How a firm resolves this trade-off depends on the cost of leaving the vacancy open and on the productivity loss from a poor match. To organize the empirical analysis, I outline a simple search model in which a firm screens a stream of applicants who differ in match quality. The model jointly determines vacancy duration, the match quality of hires and their wages, yielding predictions that I test in the empirical analysis.

The theoretical framework assumes that firms with an open vacancy meet applicants sequentially. Applicants arrive according to a Poisson process with rate $\lambda(\theta)$, with $\lambda' < 0$, that decreases in labor market tightness θ . A tighter labor market means that there are fewer applications per vacancy. The mass-layoff shock studied below lowers θ , raising the supply of applicants. Each applicant is characterized by match quality μ , drawn independently from a continuous distribution F with support $[\underline{\mu}, \bar{\mu}]$. A hired worker of match quality μ yields flow output $y(\mu)$, with $y' > 0$, and the match dissolves

at rate $\delta(\mu)$, with $\delta' < 0$. Expected tenure is thus given by $T(\mu) = 1/\delta(\mu)$, which increases in match quality μ . The firm discounts the future at rate r and pays a flow cost $c > 0$ when the vacancy remains unfilled.

The firm's only decision is which applicants to accept. Let $J(\mu)$ be the value to the firm of filling the vacancy with an applicant of quality μ , and V the value of keeping it open. The value of an open vacancy satisfies

$$rV = -c + \lambda(\theta) \int_R^{\bar{\mu}} [J(\mu) - V] dF(\mu). \quad (1)$$

Because J is increasing in match quality, the optimal policy is a reservation rule: the firm hires the first applicant whose quality exceeds a threshold R , and continues searching otherwise. The hazard rate of vacancy filling is then given by $\lambda(\theta) [1 - F(R)]$.

When the firm and an accepted applicant ($\mu \geq R$) form a match, they bargain over its surplus. If they fail to agree, the firm's outside option is the value V of continuing to search and the worker's is the value U of remaining unemployed. With worker bargaining power β , the wage divides the surplus:

$$w(\mu) = U + \beta (y(\mu) - U - V). \quad (2)$$

The continuation value V is the same object that governs selectivity above, which is what ties wages to the other outcomes. The wage responds to match output and to the firm's outside option through two opposing effects:

$$\frac{\partial w}{\partial y} = \beta > 0 \quad \text{and} \quad \frac{\partial w}{\partial V} = -\beta < 0. \quad (3)$$

The first (positive) effect arises because better-matched workers are more productive, which increases the surplus to bargain over. The second (negative) effect arises because workers have a weaker bargaining position if the value of keeping the vacancy open increases.

Using this model, I next analyze the impact of a positive supply shock. Consider a decrease in labor market tightness, which raises the arrival rate λ . In this situation, the option of holding out for a better applicant gains value, so V rises and, through $J(R) = V$, the reservation quality R rises as well. Consequently, average match quality

$\mathbb{E}[\mu \mid \mu \geq R]$, and hence job tenure, increases as the market loosens.

Instead, the impact on vacancy duration is less clear. A lower θ raises the arrival rate λ and the filling hazard directly, while the higher reservation bar R lowers it. Two mechanisms thus pull in opposite directions: a larger pool delivers more applicants per period, which accelerates filling, but it also makes the firm choosier, which slows filling. The first mechanism prevails as long as the firm does not become overly selective in response.¹ This is the case when leaving a vacancy open is relatively costly, so the firm cannot afford to wait indefinitely. Under this condition, a looser market shortens vacancy duration even as firms become more selective. Abundant applicants thus let firms fill faster and match better at the same time.

The net effect on the wage,

$$\frac{dw}{d\theta} = \underbrace{\beta \frac{d\bar{y}}{d\theta}}_{<0 \text{ (productivity)}} - \underbrace{\beta \frac{dV}{d\theta}}_{>0 \text{ (outside option)}}, \quad (4)$$

is ambiguous as well. A positive supply shock, a fall in θ , raises wages through better matches but lowers them through improved employer outside options. As a result, better matches need not be better paid. The wage rises when the match-quality channel dominates and falls when the outside-option channel dominates.

The same continuation value links an increase in applicants to worker sorting. Suppose the firm obtains a (taste-based) premium from hiring its incumbent worker type, applying a stricter reservation quality to other applicants. When the market is loose, the option cost of declining an atypical applicant is low, so the firm more frequently refrains from hiring against its preference. When applicants are scarce instead, the cost of continued search forces it to hire across groups. A fall in tightness therefore widens the hiring gap between in-group and out-group applicants, reinforcing firm segregation. I test this implication by estimating how the supply of applicants affects hiring by gender and nationality across firms with varying workforce composition.

¹The filling hazard $\lambda(\theta) [1 - F(R)]$ rises as θ falls whenever the acceptance probability $1 - F(R)$ declines less than one-for-one with the arrival-rate increase.

3 Setting and data

3.1 Vacancies

The empirical analysis combines data on vacancies, applications and employment spells in the Austrian labor market. The vacancy database covers the universe of job postings reported to the Austrian public employment service (AMS), which hosts the largest job board in Austria. In recent years, the job board has covered about 50–60 percent of all vacancies in Austria.² Unlike most private job boards, the public employment service actively manages its vacancy database and stays in contact with advertising employers. Most importantly, it regularly checks back with employers whether and when vacancies are filled. The database records a vacancy’s posting date, its expected job start date and its removal date. I define vacancy duration as the days between posting and removal, and control for the days between posting and the expected job start. The database also links each vacancy to the respective firm and occupation, using a detailed 6-digit occupation classification.

3.2 Hired workers

To examine hiring outcomes, I link the vacancies to records from the Austrian social security database.³ In addition to basic worker characteristics, the database records all employment spells subject to social security contributions as well as the corresponding earnings, and covers firm characteristics including the sector at the 4-digit NACE level. From these data, I construct the hired workers’ job duration (in days) and the average daily wage in the first year. To study worker sorting across firms, I additionally use the gender and nationality of hired workers.

3.3 Applications

Active vacancy listings are accessible through the AMS online job board, which does not require prior registration. Moreover, AMS caseworkers, who support the job search

²Statistics Austria conducts an annual vacancy survey at the firm level, which allows the AMS coverage rate to be backed out.

³Zweimüller et al. (2009) provide a comprehensive description of the data source.

of unemployed workers, use the database to refer them to suitable vacancies. When jobseekers register as unemployed, the assigned caseworker typically schedules an in-person meeting within the first 10 days. Screening the worker profiles and the stock of vacancies, caseworkers ask jobseekers to apply to suitable vacancies.⁴ In the analysis, I exploit these applications as a measure of demand for individual vacancies. Specifically, I use the number of applications during a vacancy’s first week of posting as a proxy for the vacancy-specific supply of applicants. Restricting attention to the first week captures the supply faced by the employer early in the search process, when most hiring decisions are taken, and avoids mechanical correlation with how long a vacancy happens to stay open.

3.4 Estimation sample

The estimation sample consists of filled vacancies for regular, permanent employment posted on the AMS job board in the years 2002 to 2019. As in Halla et al. (2020), I exclude sectors with strong seasonality in hiring and layoffs (hospitality, construction and agriculture) to avoid confounding seasonal employment changes with mass layoffs. Hiring outcomes are linked via firm identifiers connecting the vacancy database to the social security records, available for a subset of firms. To match workers to filled vacancies, I assign the closest firm hire within a 90-day window around the date on which a vacancy is removed. If several workers are hired on the closest date, one is chosen at random. With this procedure, I can match workers for about half of all job postings, yielding an estimation sample of around 1.2 million vacancies with observed hiring outcomes.

Table 1 provides descriptive statistics for the matched vacancy-worker sample. Vacancies receive on average about 5.6 applications in the first week, with large differences across postings. Employers need on average 51 days to fill a vacancy, and the mean time between posting and expected job start is about two to three weeks. Figure A.1 in the Appendix plots the relative frequencies of filling duration in weeks: half of all vacancies

⁴Whether a job is deemed suitable depends on a set of formal rules, including the job’s education and physical requirements, the pay difference to the prior job, the distance to work and the working hours. The public employment service describes the criteria on its website: www.ams.at/arbetsuchende (in German).

are filled within about five weeks, and around 15 percent take longer than three months. The average job lasts about 14 months and pays 68 euros per day. About 46 percent of hires are female and 77 percent are Austrian nationals. Most jobs require either compulsory education or vocational training, which reflects that the job board of the public employment service over-represents vacancies for lower- and medium-educated workers. The hiring firms are relatively sizable on average, with more than 40 percent employing over 50 workers, and their workforce is on average 44 percent female and 21 percent foreign.

4 Empirical strategy

As outlined in Section 2, changes in the supply of applicants can affect vacancy duration as well as job tenure and wages of hired workers. To quantify these effects, I estimate for each outcome y_i

$$y_i = \alpha + \beta \text{appl}_i + \gamma d_i + \psi_k + \kappa_o + \lambda_t + \varepsilon_i, \quad (5)$$

where appl_i is the number of applications to vacancy i in its first week of posting. ψ_k , κ_o , λ_t are firm, occupation and posting-month fixed effects. The variable d_i controls for the days between posting and the expected job start. Firm fixed effects absorb differences in hiring practices and persistent differences across sectors and regional labor markets. The granular occupation classification removes occupation-specific correlations between the extent of applications and the search process, and the posting-month indicators absorb seasonality and longer-run trends.

Because the number of applications may still be correlated with unobserved determinants of vacancy duration and other hiring outcomes, I instrument it with variation caused by mass layoffs.⁵ When a firm sheds a large share of its workforce, other firms operating in the same sector and location should experience a surge in applications for their vacancies. I measure this inflow at the level of the firm's sector (4-digit NACE) and district, summing over the three months before the vacancy is posted. To capture

⁵I follow the worker-flow correction of Fink et al. (2010) to distinguish genuine layoffs from spin-offs and takeovers in the Austrian worker-firm data.

Table 1: Descriptive statistics

	Mean	Std. Dev.	Median
<i>Applications</i>			
# applications	5.63	6.87	3.00
Any application	0.83	0.38	1.00
<i>Outcomes</i>			
Vacancy duration (days)	51.01	46.59	36.00
Daily wage (euros)	68.28	33.28	63.63
Job duration (days)	410.56	418.84	212.00
Female hire	0.46	0.50	0.00
Foreign hire	0.23	0.42	0.00
<i>Vacancy characteristics</i>			
Days to expected job start	16.51	30.30	4.00
Full-time	0.76	0.43	1.00
White collar	0.53	0.50	1.00
<i>Required education</i>			
Compulsory	0.29	0.45	0.00
Vocational	0.49	0.50	0.00
Higher voc.-tech.	0.15	0.35	0.00
University	0.08	0.27	0.00
<i>Firm characteristics</i>			
Foreign-worker share	0.21	0.21	0.16
Female-worker share	0.44	0.31	0.40
Avg. coworker age	38.05	5.41	38.15
Firm size ≤ 5	0.17	0.38	0.00
Firm size 6–20	0.23	0.42	0.00
Firm size 21–50	0.17	0.37	0.00
Firm size > 50	0.43	0.49	0.00

Note: The table shows summary statistics for the estimation sample, which includes 1,199,560 vacancies.

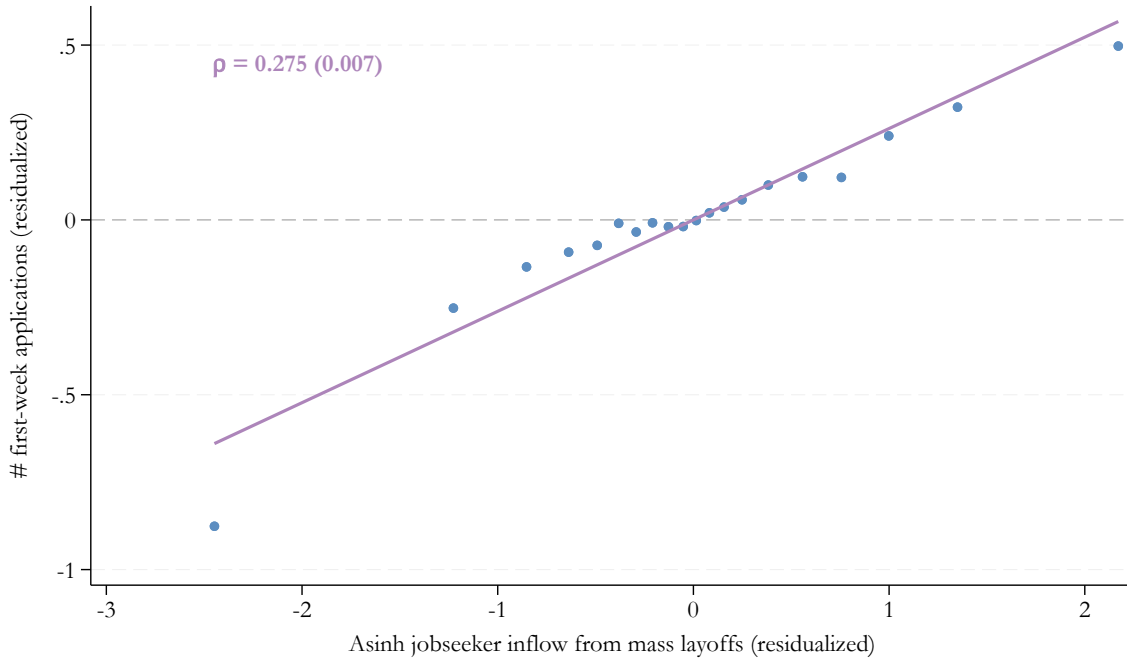
log changes while accommodating zero values, I use the inverse hyperbolic sine of the inflow (Card et al., 2022). For the construction of the instrument, I count all jobseekers who previously worked at firms losing at least 20 percent of their workforce in a single month. The conventional threshold in the displacement literature is a 30-percent

employment drop, often measured over a year (Jacobson et al., 1993; Davis and von Wachter, 2011). Because I measure the drop over a single month, a 20-percent monthly fall is a comparably sharp event. Moreover, the 20-percent threshold delivers a stronger first stage. I report the full grid of thresholds in Section 6 and show that the results are stable across these choices, with the 30-percent threshold yielding, if anything, larger effects. Because the instrument varies at the level of a sector \times district \times month cell, I cluster standard errors at this level, which allows for arbitrary correlation among vacancies exposed to the same local shock.

The instrument is a strong and stable predictor of the number of applications. Figure 1 shows the first stage as a binned scatter plot of applications against the instrument, both residualized on the controls and fixed effects of Equation (5). The relationship is strongly positive and close to linear, with a first-stage coefficient of 0.275 and an F -statistic of about 1,400. Because the instrument enters in inverse-hyperbolic-sine form, the coefficient is approximately a semi-elasticity: a doubling of the local job-seeker inflow from mass layoffs raises the number of first-week applications by about $0.275 \times \ln 2 \approx 0.19$, roughly three to four percent of the mean of 5.6. The effect is modest in level but very precisely estimated.

The identifying assumption is that, conditional on the fixed effects, mass layoffs affect hiring outcomes only through the supply of applicants. This assumption is close to that of a standard instrumental-variables (IV) design but differs in the scope of the endogenous variable. I do not assume that mass layoffs affect hiring outcomes only through the applications that I observe. The exclusion restriction is instead that they operate only through the supply of applicants, and I treat the number of applications as a representative measure of that supply, in the sense that the instrument’s effect on applications is proportional to its effect on the (unobserved) total number of applications. Under this assumption, the IV estimates identify the causal effect of applicant supply, expressed in units of observed applications: one observed application corresponds to a proportional increase in total applications, so the coefficient can be converted into the effect per total application by multiplying it by the observed share (about 0.18, see Section 5.2). Effects per standard deviation of applications are invariant to this choice of units.

Figure 1: Applications and jobseeker inflow from mass layoffs (First stage)



Note: The graph shows a binned scatter plot of the number of first-week applications against the instrument (the asinh jobseeker inflow from sector \times district mass layoffs in the three months before posting). Both variables are residualized using the set of covariates in Equation (5). The dots are means within twenty equally sized bins of the residualized instrument. The line represents the first-stage slope ρ , which is also shown in the upper left corner (with the standard error in parentheses). The estimation sample includes 1,199,560 vacancies. Standard errors are clustered at the sector \times district \times month level.

As direct support for its validity, I show next that the instrument is largely uncorrelated with predetermined firm and vacancy characteristics. Table 2 regresses a broad set of characteristics on the instrument, within the fixed-effect structure of Equation (5). These variables cover the firm's worker composition (foreign-worker and female-worker shares, average coworker age) and size, and the vacancy's working time, job type and required education as well as the days to the expected job start. The estimated standardized associations are in most cases economically negligible. A one standard-deviation change in the instrument moves each characteristic by at most about half a percent of its own standard deviation, and most by far less. Some estimates are nonetheless statistically significant, and the joint test rejects with $F = 5.7$, but only because of

the very large sample. With over a million vacancies, even very small differences are estimated precisely enough to reject.

Table 2: Balancing tests

	Std. effect	<i>t</i> -stat		Std. effect	<i>t</i> -stat
<i>Firm composition</i>			<i>Vacancy</i>		
Foreign-worker share	-0.001	-0.69	Full-time	0.002*	1.89
Female-worker share	0.005***	6.76	Part-time	-0.003**	-2.57
Avg. coworker age	0.001	0.88	White collar	-0.001	-1.18
<i>Firm size</i>			Days to exp. job start	0.001	0.50
≤5 employees	0.003**	2.53	<i>Required education</i>		
6–20	0.001	0.39	Compulsory	0.000	0.48
21–50	-0.001	-0.38	Vocational	-0.001	-0.49
>50	-0.002	-1.43	Higher voc.-tech.	0.000	0.18
			University	0.000	0.01
Joint balancing test (<i>F</i> -statistic): 5.74					

Note: Each row reports the standardized effect of a one standard-deviation change in the instrument on a predetermined firm or vacancy characteristic, controlling for the same set of covariates as in Equation (5). The estimation sample includes 1,199,560 vacancies. Standard errors are clustered at the sector×district×month level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5 Impact on vacancy duration, job tenure and wage

5.1 Regression estimates

Table 3 reports the IV estimates of Equation (5) for the three main hiring outcomes: vacancy duration, the job tenure of hired workers and their starting wages. A larger number of applications significantly reduces vacancy filling time: one additional application lowers vacancy duration by about 0.9 percent. The magnitude is modest but economically relevant: a one standard-deviation increase in applications cuts about three days from the average 51-day vacancy.

To see at which stage of the search process filling improves, Figure 2 plots IV estimates of the effect of one application on the probability that a vacancy is still unfilled in each week after posting. The effect is concentrated in weeks five to eight, when about half of the vacancies still need to be filled. In this period, one additional

Table 3: Application effect on main hiring outcomes

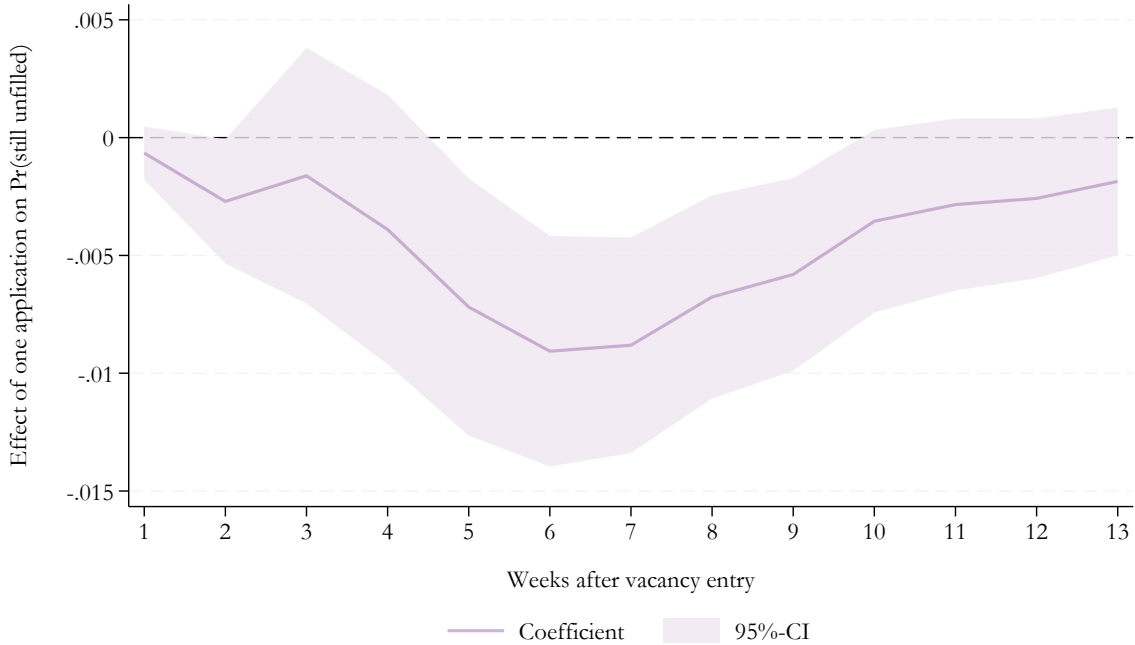
	log(vacancy duration)	log(starting wage)	log(job duration)
# applications	-0.90** (0.38)	-0.45** (0.20)	2.14*** (0.78)
Mean of dep. var.	51.01	68.28	410.56
SD of dep. var.	46.59	33.28	418.84

Note: The table shows IV estimates of Equation (5). The regressor is the number of first-week applications, instrumented with the asinh of registered jobseekers from mass layoffs in the same sector and district summed over the three months before posting. The first stage, common to the three outcomes, has a coefficient of 0.275*** (standard error 0.007) and an F -statistic of 1400. The regression coefficients are multiplied by 100. The mean and standard deviation of the dependent variable are reported in levels (vacancy and job duration in days, wage in euros per day). The estimation sample includes 1,199,560 vacancies. Standard errors are clustered at the sector \times district \times month level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

application reduces the probability that a vacancy is still active by one percentage point.

What is the effect of additional applications on job tenure? If a larger supply of applicants allows employers to find better matches, early separations should become less frequent. As shown in Table 3, more applications indeed significantly raise job tenure: one additional application increases job duration by about 2.1 percent. Yet, longer tenure need not signal a better match. A worker with poorer outside options also stays longer, simply because there is less to leave for. This is a genuine concern here, because the instrument covers displaced workers, whose outside options are weak, so they might stay in a job even when the fit is poor. To better distinguish between the two mechanisms, I next examine how the effect evolves over the spell rather than its impact on total tenure alone. Because a poor match tends to reveal itself quickly, separations due to mismatches are expected to be concentrated in the first months after hiring. A weaker outside option, by contrast, lowers the separation hazard at every horizon, so it would raise survival by a margin that does not fade, and if anything widens, as the spell lengthens. Figure 3 shows the first pattern and not the second: the effect on survival is largest in the first months after hiring and does not widen as the spell lengthens. A deeper applicant pool therefore appears to lengthen tenure by improving the match, not merely by hiring workers who have fewer outside options. Together

Figure 2: Application effect on vacancy duration

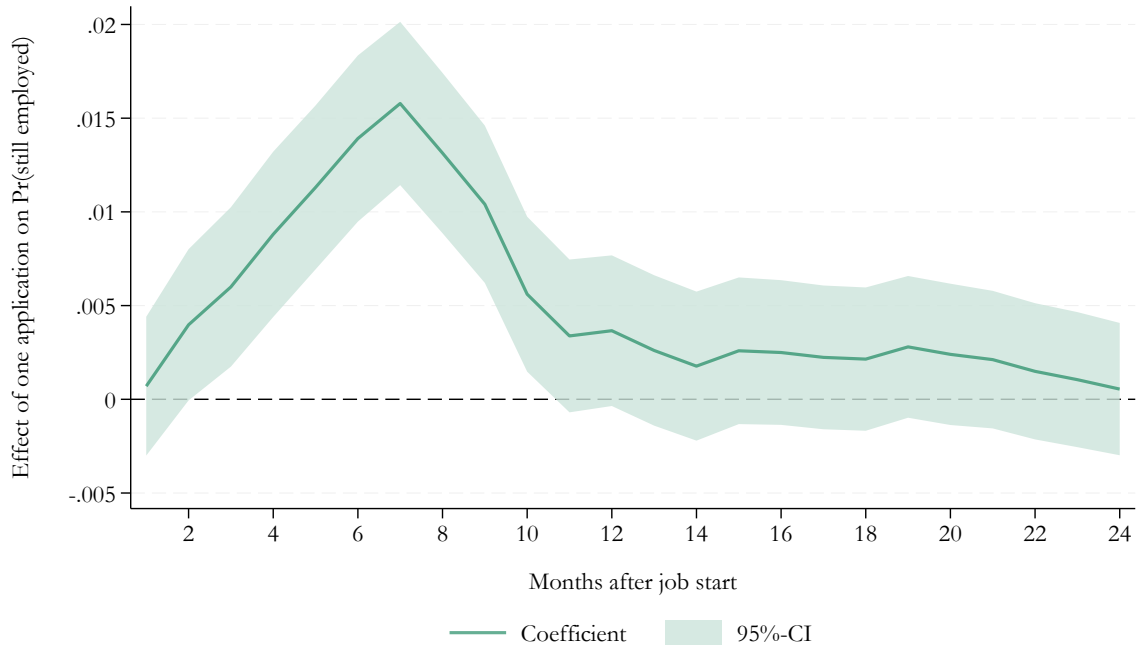


Note: The graph shows IV estimates of Equation (5) where the outcome is an indicator for the vacancy still being unfilled in week t after posting. The shaded area shows the 95% confidence interval. The estimation sample includes 1,199,560 vacancies. Standard errors are clustered at the sector \times district \times month level.

with the duration result, these estimates are consistent with the model assuming that keeping vacancies unfilled is relatively costly: when many jobseekers apply, firms fill faster and select better-matching workers.

As discussed in Section 2, a larger applicant pool may affect wages in two opposite directions. By making the firm more selective, it raises the productivity of the hire, which should push pay up. But the increase in applicants strengthens the firm’s wage-setting power, since a larger pool improves its outside option, lets the firm keep a larger share of the match surplus and thus pushes wages down. The net effect is therefore ambiguous a priori. As shown in Table 3, more applications significantly reduce starting wages, by about 0.45 percent per application. Together with the positive effect on job tenure, this suggests that a better-matched worker is hired at a lower wage when applicants are abundant. The negative effect due to stronger bargaining power of

Figure 3: Application effect on job tenure of hired workers



Note: The graph shows IV estimates of Equation (5) where the outcome is an indicator for the hire still being employed in month t after job start. The shaded area shows the 95% confidence interval. The estimation sample includes 1,199,560 vacancies. Standard errors are clustered at the sector \times district \times month level.

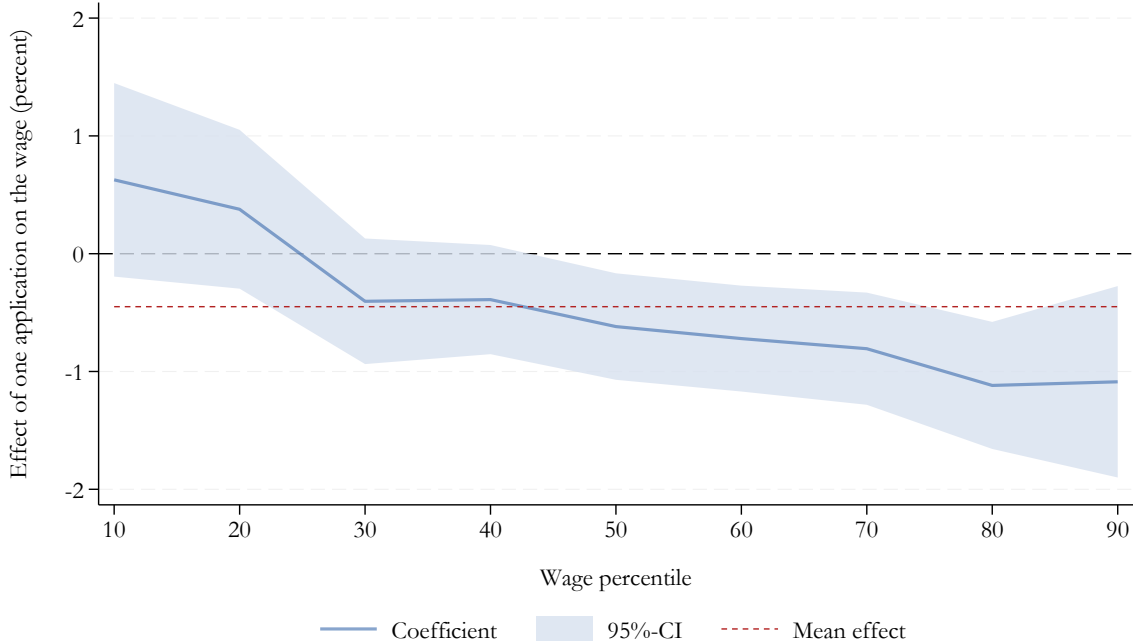
employers outweighs the positive effect due to productivity gains with better-matching workers.

This average decline conceals a clear gradient across the wage distribution. Figure 4 plots the application effect around each decile of the starting wage, estimated by instrumenting the application count in a recentered-influence-function (RIF) regression so that each coefficient is the effect on that decile (Firpo et al., 2009). The response is close to zero at the bottom and grows steadily toward the top, where it reaches about one percent per application above the median. A deeper applicant pool thus lowers pay where the wage is discretionary and has little effect at the bottom, where collective bargaining floors bind and leave little room to adjust downward.⁶ This is the pattern that employer wage-setting power implies: the supply shock affects the wages a firm is

⁶Because Austria has nearly universal collective bargaining coverage, the wage floors set by collective agreements cover essentially every hire.

free to set, which lie higher up the distribution where no floor binds.

Figure 4: Application effect on wages of hired workers by wage decile



Note: The graph shows IV estimates of Equation (5) where the dependent variable is the recentered influence function (RIF) of the log wage at the respective decile (Firpo et al., 2009). Each estimate is the decile-specific effect of one additional application on the wage (in percent). For quantile τ , with cutoff q_τ and density of log wages $f(q_\tau)$ at that cutoff, each hire's log wage w is recoded as $q_\tau + (\tau - \mathbf{1}\{w \leq q_\tau\})/f(q_\tau)$. The density is estimated with a Gaussian kernel. The dashed line indicates the mean effect. The shaded area shows the 95% confidence interval. The estimation sample includes 1,199,560 vacancies. Standard errors are clustered at the sector \times district \times month level.

5.2 Interpretation of effect sizes

The presented IV coefficients show the impact per one additional application. The magnitudes of these effects are more easily interpreted when scaled to a typical change in applicant supply. A one standard-deviation increase (about 6.9 applications) shortens vacancy duration by 6.1 percent, raises job tenure by 14.7 percent and lowers starting wages by 3.1 percent. Because I use an IV design to estimate the effects of changes in applicant supply, the magnitudes should be interpreted as local effects. The instrument identifies the impact for compliers, i.e. the vacancies whose applicant pool grows with

nearby mass layoffs, and it does so through the marginal applicants those layoffs supply. The estimates therefore reflect the effect of one more applicant of this kind at a given applicant pool size.

Because the applications used in this analysis stem from job referrals at the public employment service, they capture only a subset of all applications. Assuming that referred and outside applications are equally likely to fill a vacancy, I can back out the expected number of total applications. Since the share of vacancies filled by a referred jobseeker is about 18 percent, scaling up the referral-based applications implies about 31 applications per vacancy on average in the first week after vacancy entry. Interpreting the IV estimates as the effect of a unit change in overall applications would be equivalent to multiplying the reported IV estimates by 0.18. Importantly, the per-standard-deviation effects above are invariant to this scaling, since the coefficient and the standard deviation of applications scale by offsetting factors, leaving their product unchanged.

The 3.1 percent decline in the starting wage for a standard deviation rise in applications is sizable. It is comparable to one or two years of average real wage growth, and it implies that a larger applicant pool moves a non-trivial share of the match surplus from the worker to the firm. The real-side effects are equally concrete. The same one-standard-deviation increase shortens the average vacancy by about three days, from a base of 51, and lengthens the average hire's tenure by about two months, from a base of 14 months, a clear improvement in match quality and matching efficiency. Taken together, the three effects describe a single shift rather than three unrelated ones. A deeper applicant pool allows a firm to fill its vacancy faster, hold out for a better-matching worker and pay that worker less. The same forces that improve the match also strengthen the firm's position in setting the terms of employment.

6 Robustness

For the analysis, I assume that inflows from mass layoffs represent a genuine supply shock that affects hiring only through the size of the applicant pool. To provide further evidence on the validity of this assumption, I run two robustness checks for the

instrument.

If the instrument captures a time-specific supply shock rather than persistent local conditions, only layoffs that occur before a vacancy is posted should affect its outcomes. Panel A of Table 4 reports reduced-form regressions of each outcome on the layoffs in a given window. The pre-posting window moves all three outcomes, whereas layoffs that fall 7–9 and 10–12 months after posting, in months after the typical vacancy is filled, have no significant effect on wages or vacancy duration, and only a marginally significant effect on job tenure in the 7–9 month window. The instrument thus reflects a surge of applicants around the time of posting rather than a lasting feature of the local market, which would have moved outcomes in the later windows as well.

Table 4: Placebo tests and layoff definitions

	<i>F</i> -stat	log(vac. dur.)	log(wage)	log(job dur.)
Panel A. <i>Placebo: Post-filling layoffs (reduced form)</i>				
Pre-posting $t - 1$ to $t - 3$		-0.25** (0.11)	-0.13** (0.06)	0.60*** (0.22)
Post-filling $t + 7$ to $t + 9$		-0.11 (0.11)	0.00 (0.06)	0.37* (0.22)
Post-filling $t + 10$ to $t + 12$		-0.10 (0.11)	-0.08 (0.06)	-0.03 (0.23)
Panel B. <i>Layoff definition: Jobseekers $\geq x\%$ (IV estimates)</i>				
$\geq 10\%$	1527	-0.93***	-0.14	2.12***
$\geq 20\%$	1400	-0.90**	-0.45**	2.14***
$\geq 30\%$	1003	-1.50***	-0.61***	3.38***
$\geq 40\%$	617	-1.27**	-0.74***	4.70***
$\geq 50\%$	381	-1.02	-0.86***	5.13***

Note: Panel A reports reduced-form regressions of each outcome on the asinh number of laid-off jobseekers in the indicated window relative to vacancy-posting month t . Panel B reports, for each threshold, the first-stage F -statistic and the IV coefficients, where the instrument counts registered jobseekers from firms losing at least the indicated share of their workforce in one month. All coefficients are multiplied by 100. The estimation sample includes 1,199,560 vacancies. Standard errors are clustered at the sector \times district \times month level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

As another robustness check, I vary the definition of the layoff measure. Panel B of Table 4 reports the first-stage strength and the three IV estimates across the grid of jobseeker layoff thresholds, from 10 to 50 percent of a firm’s workforce. The first stage is stronger at the lower thresholds, while the IV estimates keep the same sign across the whole grid and stay significant at all but the most extreme cutoffs. If anything, the IV estimates point to somewhat larger effects at higher thresholds, which is consistent with the local nature of the instrument: a bigger mass layoff is a more salient supply shock, and the marginal applicant it sends to a competing vacancy need not fully correspond to that of a smaller one.

7 Impact on worker segregation

7.1 The segregation gradient

When the supply of jobseekers influences the hiring patterns shown above, changes in tightness should also affect the composition of workers across firms. Many workplaces are segregated, for instance, along gender or nationality, driven by both supply and demand factors. When labor markets are loose, employers can be more selective and may prefer to hire workers similar in characteristics to their existing workforce. This should be the case if the current workforce composition reflects employer preferences that are stable over time. A lower number of applicants forces firms instead to also consider workers from less represented groups. To test this, I relate the number of applications to the gender and nationality of hires and examine how the relation varies with the firm’s existing composition. As a first step, I sort firms into groups based on the share of female (foreign) workers among their existing workforce. Second, using an indicator for hiring a female (foreign) worker as the outcome, I estimate IV regressions that interact the instrumented number of applications with the group dummies and use the same controls shown in Equation (5). Figure 5 plots the resulting group-specific application effects. The upper panel shows that firms with few female workers become less likely to hire a woman when they receive more applications, whereas firms with many female workers become more likely to do so. The coefficient rises with the firm’s female share, though the pattern is not perfectly monotonic. The lower panel shows a

similar gradient for nationality, where more applications reduce the likelihood of hiring a foreign worker in firms with few foreigners and raise it in firms with more. For this outcome, the high-share groups are estimated imprecisely because few firms employ a very high share of foreign workers.⁷ Both gradients indicate that a larger applicant pool reinforces the firm’s existing composition. In the next subsection, I examine whether this works through changes in the composition of the applicant pool or through the firm’s own selection.


To summarize each gradient in a single coefficient and test its significance, I interact the instrumented number of applications with the firm’s pre-hire share of female (foreign) workers, which directly tests whether the effect varies with that share. Table 5 reports these estimates. For gender, the effect of applications on hiring a woman is negative at firms with no female workers and rises significantly with the firm’s female share: the estimated slope is 3.1 percentage points. This shows again that more applicants make male-dominated firms hire even fewer women. The second, quadratic specification in the next column shows that the relationship is concave. For nationality, the gradient is also positive and significant, with a slope of 5.4 percentage points. Here the estimates of the quadratic specification do not indicate a non-linear relation.

These results are consistent with the hypothesis that tighter labor markets attenuate firm segregation. When job applicants are scarce, employers widen their search and also hire atypical workers. This connects to the classic argument of Becker (1957) that competition erodes the scope for discrimination: with few applicants, the cost of taste-based hiring rises, so firms hire more broadly. Adding to similar findings in the literature (Biddle and Hamermesh, 2013; Weber and Zulehner, 2014), I use the vacancy-level supply shock to isolate the same discrimination mechanism at the level of the individual hire. The estimated differences are economically relevant. A one standard-deviation increase in applications widens the gap in the probability of a female hire between a firm in the upper quartile of the female-worker share and one in the lower quartile by roughly 10 percentage points. For comparison, the average female-hire rate is 46 percent (see Table 1). The nationality gradient is of similar size, about 10


⁷To get a sufficiently large sample, I use a broader definition for the foreigner share, grouping together all firms with at least 70 percent of non-native workers.

Figure 5: Application effect on worker segregation

(a) Impact on hiring a female worker



(b) Impact on hiring a foreign worker



Note: The graph shows IV estimates for the indicated group of the firm's pre-hire female (foreign) worker share, where the outcome is an indicator for a female (foreign) hire. The estimates are obtained from a regression that interacts the instrumented application count with the group dummies and controls for the same set of covariates as in Equation (5). The top foreign-share deciles are pooled into the (0.7, 1] group. The bars show 95% confidence intervals. The estimation sample includes 1,199,560 vacancies. Standard errors are clustered at the sector \times district \times month level.

percentage points across the same quartiles, relative to a 23 percent foreign-hire rate.

Table 5: Application effect on hiring a female (foreign) worker

	Female hire		Foreign hire	
	Linear	Quadratic	Linear	Quadratic
# applications	-0.61** (0.24)	-1.34*** (0.29)	-1.66*** (0.29)	-1.61*** (0.34)
# applications \times share	3.13*** (0.49)	9.09*** (1.74)	5.39*** (0.71)	5.46*** (1.88)
# applications \times share ²	–	-6.75*** (1.75)	–	-0.49 (2.32)
Share	-17.32*** (2.92)	-45.56*** (10.32)	-21.02*** (4.21)	-30.03*** (10.89)
Share ²	–	32.56*** (10.44)	–	14.52 (13.77)

Note: The table shows IV estimates, where the outcome is an indicator for a female (foreign) hire. The estimates are obtained from a regression that interacts the instrumented application count with the firms' female (foreign) worker shares and controls for the same set of covariates as in Equation (5). Coefficients are multiplied by 100. The estimation sample includes 1,199,560 vacancies. Standard errors are clustered at the sector \times district \times month level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

7.2 Supply composition or employer selection?

The segregation pattern documented above has two possible sources. A larger applicant pool could raise own-type hiring mechanically, because the pool a firm faces is itself segregated: occupational sorting channels female and foreign jobseekers toward firms that already employ them, so more applicants can simply mean more own-type applicants. Alternatively, a surge in applicants could allow employers to act on a preference for similar workers and hire more own-type workers than the pool's composition would imply. The first is a supply-composition channel, while the second reflects employer selection. I separate them using the gender and nationality of the applicants themselves.

First, I replicate the segregation analysis from the previous subsection but focus on changes in the composition of applications rather than on the hiring of female (foreign)

workers. Specifically, I define as outcome the female (foreign) share among applicants.⁸ Figure A.2 in the Appendix shows the corresponding estimates. Compared to the results for the hiring regression shown in Figure 5, effect sizes are overall smaller in both panels. For gender, the coefficients are negative throughout and grow more negative at firms with a higher female share. This implies that more applications make the pool slightly less female rather than more. For nationality, the effect is positive for all firm groups and rises with the firm’s foreign share.

To separate employer preferences from changes in the applicant pool, I contrast these results with the previous estimates for hiring. Both the hire and the pool are own-type shares on the same scale. The hire is the own-type share of the one worker hired, and the pool is the own-type share of the applications. If a firm hired a random applicant from its pool, the two shares would respond identically to the supply shock. I therefore estimate the same interaction as before, applications times the firm’s own-type share, on three outcomes: the own-type hire indicator (*Total*), the own-type share of applications (*Supply*) and their difference (*Selection*). Because the third outcome is the first minus the second, the gradients add up. *Supply* is the gradient that random hiring from the pool would produce, and *Selection* is the part due to employer choice.

Table 6 shows that the effects on the two margins differ strongly. For gender, the supply component is significantly negative: the shock does not make the pool more female at firms with a high female share but slightly less. Thus, changes in gender-specific hiring patterns can be attributed to employer selection. For nationality, the supply component is positive and accounts for about 40 percent of the gradient. Yet, the remaining selection effect is still large. This shows that, in both cases, employer selection is the dominant cause of changes in the hiring of female (foreign) workers. For this decomposition analysis, I have to assume again that the applications I observe in the first week are a representative sample of all applications. If unobserved applicants are representative of the pool, the comparison is unaffected. If they are instead systematically more own-type than the observed applications, the estimate for *Selection* would overstate the true impact of employers.

⁸For the analysis in this subsection, I need to restrict the estimation sample to vacancies that receive at least one application in the first week, which removes about 17 percent of vacancies. Otherwise, the female (foreign) share of applicants could not be defined.

Table 6: Worker segregation: supply composition versus employer selection

	Female	Foreign
<i>Total</i>		
Outcome: Actual hire	3.39*** (0.51)	6.88*** (0.85)
<i>Supply</i>		
Outcome: Application share	-1.72*** (0.25)	2.86*** (0.44)
<i>Selection</i>		
Outcome: Actual hire - application share	5.12*** (0.57)	4.02*** (0.86)

Note: The table shows IV estimates for three different outcomes: an indicator for a female (foreign) hire (*Total*), the vacancy's female (foreign) share of applications (*Supply*) and the difference of both outcomes (*Selection*). The estimates are obtained from a regression that interacts the instrumented application count with the firms' female (foreign) worker shares and controls for the same set of covariates as in Equation (5). Coefficients are multiplied by 100. The estimation sample includes 992,674 vacancies, the subsample of vacancies that receive any first-week applications. Standard errors are clustered at the sector×district×month level. * p<0.10, ** p<0.05, *** p<0.01.

8 Conclusion

Whether a firm faces many applicants or few is usually treated as a question of recruiting cost because firms receiving fewer applications are expected to search longer. But the size of the applicant pool may also influence who gets hired eventually. It can change which worker is chosen, what that worker is paid and how the composition of the workforce shifts. These are consequences that can benefit or hurt both workers and firms.

In this paper, I study how the supply of job applicants shapes firms' hiring outcomes, using a vacancy-level measure of applicant supply built from a subset of job applications and instrumented by the inflow of jobseekers from nearby mass layoffs. This analysis is feasible because the Austrian vacancy data can be linked to records on worker employment, earnings and displacement histories, which allows me to examine effects of applicant supply shocks on vacancy filling duration, worker-firm match

quality, pay and workforce composition.

The results show that an increase in applications moves hiring along several margins at once. Firms fill vacancies faster and hire workers who stay longer, consistent with better worker-firm matches. At the same time, they pay these better-matched workers less. This wage decline is concentrated in the upper part of the wage distribution, where pay is discretionary, while collective bargaining floors define pay at the bottom. A better match would, on its own, raise pay rather than lower it. A negative wage effect suggests that a markdown due to strengthened wage-setting power of employers dominates the positive productivity effect.

Moreover, I find that a larger applicant pool reinforces firm segregation by gender and nationality. When more jobseekers apply, firms hire more workers who resemble their existing workforce. A firm with mostly male or native workers becomes even less likely to hire a woman or a foreigner, while a firm that already employs many such workers becomes even more likely to do so. A decomposition of segregation effects shows that this pattern reflects employer selection rather than a shift in the composition of applicants. This is consistent with the notion that slack labor markets enable, and tight markets attenuate, discriminatory hiring.

Due to an aging workforce in many countries, the affected economies may face increasing labor shortages in the near future. Read in reverse, my results imply that a thinner applicant pool makes vacancies harder to fill and worker-firm matches worse, but also raises wages and improves access to jobs for groups of workers that firms hire less often when applicants are abundant. This creates a trade-off for common policy responses to shortages. Measures that raise the effective supply of applicants, such as immigration or higher labor force participation, help firms fill positions and improve match quality. At the same time, greater applicant supply can depress wages and reinforce segregation across workplaces. Policies aimed at easing shortages should therefore take these distributional effects into account. In particular, interventions that improve match quality, rather than simply enlarge the applicant pool, are more likely to raise efficiency without weakening the bargaining position of workers.

References

- Abraham, K. G., Haltiwanger, J. C., and Rendell, L. E. (2020). How tight is the US labor market? *Brookings Papers on Economic Activity*, 2020(1):97–165.
- Altonji, J. G., Kahn, L. B., and Speer, J. D. (2016). Cashier or consultant? Entry labor market conditions, field of study, and career success. *Journal of Labor Economics*, 34(S1):S361–S401.
- Azar, J., Marinescu, I., Steinbaum, M., and Taska, B. (2020). Concentration in US labor markets: Evidence from online vacancy data. *Labour Economics*, 66:101886.
- Baert, S., Cockx, B., Gheyle, N., and Vandamme, C. (2015). Is there less discrimination in occupations where recruitment is difficult? *ILR Review*, 68(3):467–500.
- Becker, G. S. (1957). *The economics of discrimination*. University of Chicago Press.
- Biddle, J. E. and Hamermesh, D. S. (2013). Wage discrimination over the business cycle. *IZA Journal of Labor Policy*, 2:1–19.
- Bloom, N. and Van Reenen, J. (2011). Human resource management and productivity. In *Handbook of Labor Economics*, volume 4, pages 1697–1767. Elsevier.
- Card, D., Colella, F., and Lalive, R. (2024). Gender preferences in job vacancies and workplace gender diversity. *Review of Economic Studies*.
- Card, D., DellaVigna, S., Funk, P., and Iriberry, N. (2022). Gender differences in peer recognition by economists. *Econometrica*, 90(5):1937–1971.
- Cullen, Z., Hoffman, M., and Koenig, F. (2024). Labor shortages and firm search. Mimeo.
- Cullen, Z. and Perez-Truglia, R. (2022). How much does your boss make? The effects of salary comparisons. *Journal of Political Economy*, 130(3):766–822.
- Davis, S. J., Faberman, R. J., and Haltiwanger, J. C. (2013). The establishment-level behavior of vacancies and hiring. *Quarterly Journal of Economics*, 128(2):581–622.

- Davis, S. J. and von Wachter, T. (2011). Recessions and the costs of job loss. *Brookings Papers on Economic Activity*, 2011(2):1–72.
- Doleac, J. L. and Hansen, B. (2020). The unintended consequences of "ban the box": Statistical discrimination and employment outcomes when criminal histories are hidden. *Journal of Labor Economics*, 38(2):321–374.
- Domash, A. and Summers, L. H. (2022). How tight are US labor markets? Working Paper No. 29739, NBER.
- Fink, M., Segalla, E., Weber, A., and Zulehner, C. (2010). Extracting firm information from administrative records: The ASSD firm panel. Working Paper No. 1004, NRN.
- Firpo, S., Fortin, N. M., and Lemieux, T. (2009). Unconditional quantile regressions. *Econometrica*, 77(3):953–973.
- Halla, M., Schmieder, J., and Weber, A. (2020). Job displacement, family dynamics, and spousal labor supply. *American Economic Journal: Applied Economics*, 12(4):253–87.
- Hoeck, C. P. (2023). Wage effects of labor market tightness. Working paper.
- Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993). Earnings losses of displaced workers. *American Economic Review*, 83(4):685–709.
- Klaeui, J., Kopp, D., Lalive, R., and Siegenthaler, M. (2024). Adapting to scarcity: Job search and recruiting across occupational boundaries. Mimeo.
- Kuhn, P. and Shen, K. (2023). What happens when employers can no longer discriminate in job ads? *American Economic Review*, 113(4):1013–48.
- Manning, A. (2003). *Monopsony in motion: Imperfect competition in labor markets*. Princeton University Press.
- Mueller, A. I., Osterwalder, D., Zweimüller, J., and Kettemann, A. (2024). Vacancy durations and entry wages: Evidence from linked vacancy-employer-employee data. *Review of Economic Studies*, 91(3):1807–1841.

- Nordin, M., Persson, I., and Rooth, D.-O. (2010). Education–occupation mismatch: Is there an income penalty? *Economics of Education Review*, 29(6):1047–1059.
- Oyer, P. and Schaefer, S. (2011). Personnel economics: Hiring and incentives. In *Handbook of Labor Economics*, volume 4, pages 1769–1823. Elsevier.
- Patterson, C., Şahin, A., Topa, G., and Violante, G. L. (2016). Working hard in the wrong place: A mismatch-based explanation to the UK productivity puzzle. *European Economic Review*, 84:42–56.
- Pissarides, C. A. (2000). *Equilibrium unemployment theory*. MIT Press.
- Şahin, A., Song, J., Topa, G., and Violante, G. L. (2014). Mismatch unemployment. *American Economic Review*, 104(11):3529–3564.
- Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1):25–49.
- Turrell, A., Speigner, B., Copple, D., Djumalieva, J., and Thurgood, J. (2021). Is the UK’s productivity puzzle mostly driven by occupational mismatch? An analysis using big data on job vacancies. *Labour Economics*, 71:102013.
- Weber, A. and Zulehner, C. (2014). Competition and gender prejudice: Are discriminatory employers doomed to fail? *Journal of the European Economic Association*, 12(2):492–521.
- Zweimüller, J., Winter-Ebmer, R., Lalive, R., Kuhn, A., Wuellrich, J.-P., Ruf, O., and Büchi, S. (2009). The Austrian Social Security Database (ASSD). *University of Zurich Working Paper*, 410.

Appendix

Figure A.1: Histogram of vacancy duration


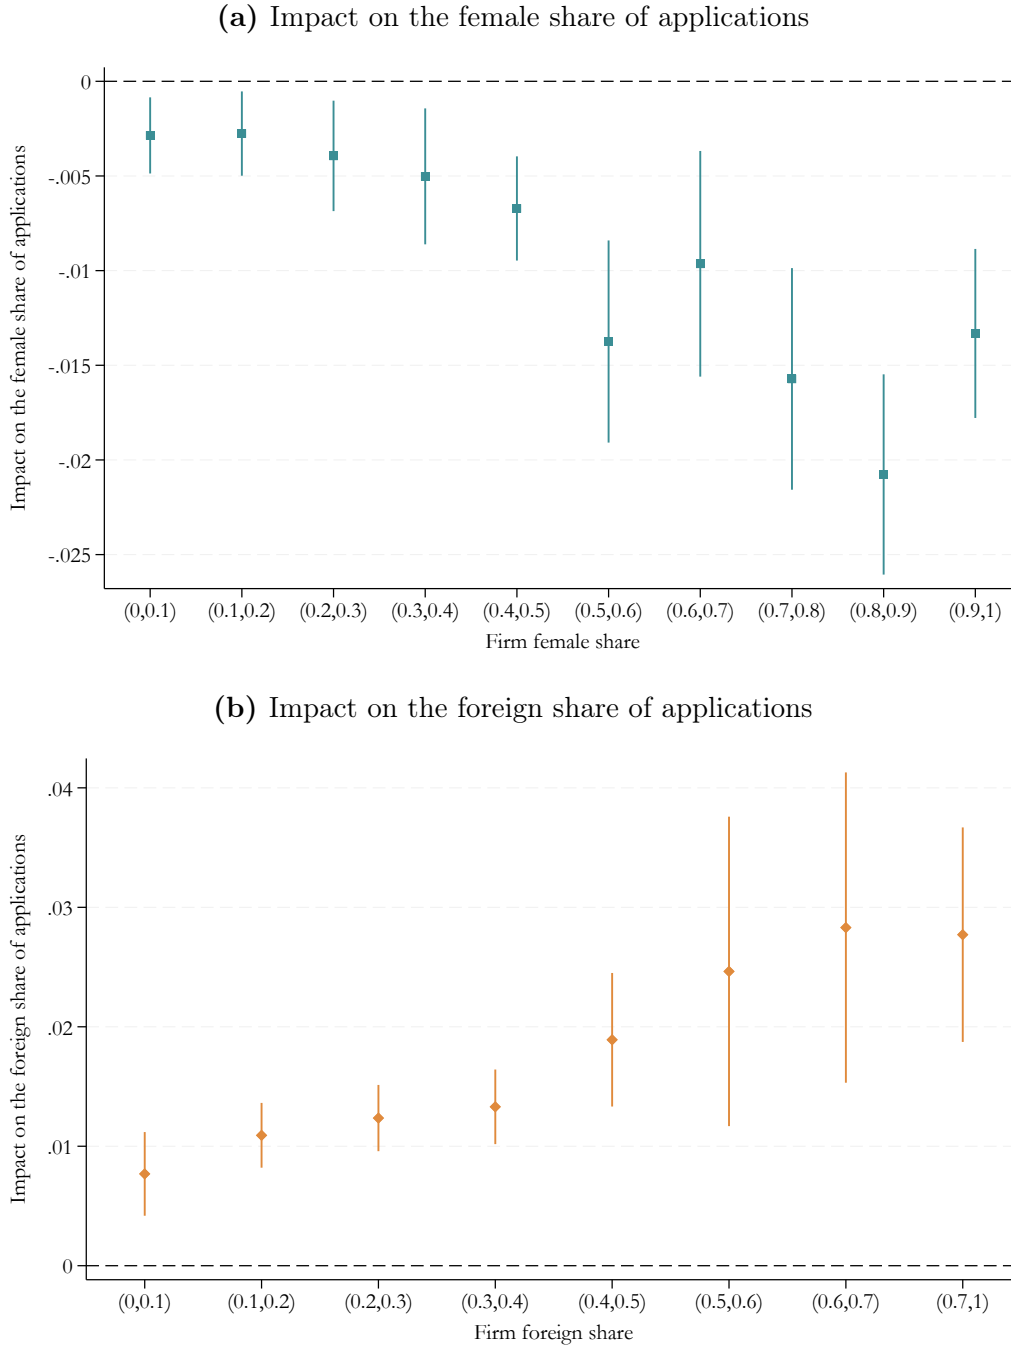


Figure A.2: Application effect on the composition of applications



Note: The graph shows IV estimates for the indicated group of the firm's pre-hire female (foreign) worker share, where the outcome is the vacancy's female (foreign) share of applications. The estimates are obtained from a regression that interacts the instrumented application count with the group dummies and controls for the same set of covariates as in Equation (5). The top foreign-share deciles are pooled into the (0.7, 1] group. The bars show 95% confidence intervals. The estimation sample includes 992,674 vacancies, the subsample of vacancies that receive any first-week applications. Standard errors are clustered at the sector×district×month level. 32