

The wage elasticity of job applications: Quantity and quality margins

Louise Demoury

Adelaide Fabbi

Roland Rathelot *

28th June 2026

Abstract

We study how an occupation-wide wage increase affects the number and quality of job applications and subsequent hiring. We use data from *Extracadabra*, an online platform for short-term hospitality work in France, which records job notifications, applications, and hiring decisions. The platform raised the default hourly wage for waiting-staff jobs from €16 to €17 in October 2023 while leaving the defaults for three neighbouring occupations unchanged. A difference-in-differences design shows that the wage increase raised applications by approximately 23%, implying a wage elasticity of about 3.5. Mean applicant quality rose by 2%, an elasticity of approximately 0.3, with the largest improvement at the bottom of the quality distribution. Hires rose by about 11%, implying an elasticity of 1.8. Because the wage change may divert applications from the comparison occupations, the usual no-interference assumption need not hold. We therefore estimate a model of application and hiring decisions that separates the equilibrium response in the treated occupation from spillovers to neighbouring occupations. The decomposition assigns about 6% of the application difference-in-differences estimate to displacement from the comparison occupations and essentially none of the quality effect to spillovers. The reduced-form estimates therefore closely approximate the equilibrium responses in the treated occupation.

Keywords: job search, hiring, online platforms, labour-supply elasticity, job applications, spillovers, minimum wages, wage floors.

JEL Codes: J22, J23, J31, J62, J63, J64, E24.

*Demoury: Institut des Politiques Publiques. Fabbi and Rathelot: Institut Polytechnique de Paris (ENSAE). We thank Extracadabra for making the data available; Felix Koenig, Rafael Lalive, Michael Oberfichtner, Barbara Petrongolo, Jon Piqueras, and Michael Siegenthaler; and participants at the 2026 Konstanz Macro-Labor Workshop for helpful comments. Guillaume Pousse provided excellent research assistance. All errors are our own. Demoury and Fabbi are also affiliated with CREST. Rathelot is also affiliated with CREST, CEPR, IZA, RFBerlin, and J-PAL.

1 Introduction

Employers in hospitality, construction, and health care often report persistent recruitment difficulties, even when unemployment remains substantial. In a frictionless labour market, wages would adjust to direct workers towards occupations with high labour demand. How strongly workers respond to such wage differences is therefore central to understanding both recruitment and the allocation of labour across occupations. Existing evidence largely concerns the number of applications or hires. Much less is known about the quality of the workers attracted by higher wages, partly because applications, applicant characteristics, and hiring decisions are rarely observed in the same data.

This article studies how an occupation-wide wage increase affects the quantity and quality of applications and subsequent hiring. We use data from *Extracadabra*, an online platform for short-term hospitality work in and around Paris. The platform sets an occupation-specific default hourly wage, which most employers adopt. On 5 October 2023, it raised the default wage for waiting-staff jobs from €16 to €17, an increase of 6.25%, while leaving the defaults for bar staff, runners, and receptionists unchanged. Comparing waiting-staff vacancies with vacancies in these three occupations before and after the change identifies the response to a change in an occupation’s relative wage rather than to a wage increase at a single firm.

The platform data record the workers notified about each vacancy, their applications, and employers’ hiring decisions. We use these hiring decisions and rich worker–job characteristics to construct an applicant-level quality index. The data therefore allow us to distinguish four margins: the number of applications, the mean quality of the applicant pool, its lower and upper tails, and realised hiring.

Our preferred difference-in-differences specification shows that the wage increase raised the number of applications by approximately 23% relative to the comparison occupations. Dividing the change in log applications by the change in the log wage gives an elasticity of about 3.5. Mean applicant quality rose by approximately 2% of its pre-treatment mean, implying an elasticity of about 0.3. The minimum quality in the applicant pool increased more strongly, with an elasticity of approximately 1.7. The estimated effect on maximum quality is small and statistically insignificant; the confidence interval rules out an elasticity above approximately 0.4. Finally, the number of hires increased by about 11%, implying an elasticity of 1.8. Thus, the wage increase attracted substantially more applicants and modestly improved the applicant pool, primarily by raising its lower tail.

The reduced-form contrast does not, by itself, identify the equilibrium effect on waiting-

staff jobs. Workers can apply across occupations, so the higher waiting-staff wage may divert applications from the comparison occupations. This interference violates the stable unit treatment value assumption underlying a conventional difference-in-differences interpretation. We therefore develop a model of workers' application decisions and employers' hiring decisions. Workers differ in their productivity in two occupations and account for both the wage and their probability of being hired. A higher wage can attract productive workers, intensify competition, and discourage applicants with weaker hiring prospects. The model separates the equilibrium response in the treated occupation from the spillover to the comparison occupations.

Estimating the model shows that spillovers account for about 6% of the model-based application contrast: the higher waiting-staff wage reduces applications to the comparison occupations, modestly amplifying the difference-in-differences estimate. The remaining 94% reflects the equilibrium response in the treated occupation. The corresponding quality spillover is negligible. The reduced-form estimates therefore provide close approximations to the equilibrium effects of the wage increase on waiting-staff vacancies.

This article contributes to three literatures. First, it contributes to work estimating the elasticity of labour supply to firms and jobs using applications, vacancy duration, recruitment, quits, and employment flows (Ransom and Sims, 2010; Staiger et al., 2010; Matsudaira, 2014; Banfi and Villena-Roldán, 2019; Dube et al., 2020; Marinescu and Wolthoff, 2020; Bassier et al., 2022; Sokolova and Sorensen, 2021; Manning, 2021; Card, 2022; Hirsch et al., 2022; Datta, 2023; Bassier et al., 2025). Banfi and Villena-Roldán (2019) and Marinescu and Wolthoff (2020) show that higher-wage jobs receive more applications. Belot et al. (2022) and Beerli et al. (2026) study how jobseekers respond to experimentally varied wage information. Hirsch et al. (2022) and Bassier et al. (2025) focus on recruitment outcomes (vacancy duration and hires). Our source of variation differs from most of the literature because the wage changes for an entire occupation. Our design allows us to capture both the response to the higher wage and the reallocation of applications across contemporaneously available jobs. To our knowledge, our paper is the first attempt to decompose this market-level effect between an effect on the affected occupation and spillovers. Another contribution is that we can measure the impact on the number of applications, application quality, and hires within the same setting.

Second, the article contributes to the smaller literature on how wages affect applicant composition. Dal Bó et al. (2013) provide experimental evidence that higher public-sector wages attract a larger and more qualified applicant pool. Related work studies how wage information and wage floors affect applicant selection, skill requirements, and the composition

of job offers (Clemens et al., 2021; Balgova et al., 2025). We observe applications, applicant characteristics, and hiring outcomes on the same platform, allowing us to distinguish changes in the number of applicants from changes throughout the quality distribution and from realised hiring.

Third, the article relates to work on wage floors, reallocation, and occupational mobility. Wage-setting policies can reallocate workers and firms across jobs, establishments, and local labour markets (Cengiz et al., 2019; Harasztosi and Lindner, 2019; Dustmann et al., 2022; Derenoncourt and Montialoux, 2021; Caldwell and Danieli, 2024; Azar et al., 2024; Derenoncourt and Weil, 2025; Giupponi et al., 2024), and can generate spillovers to untreated workers, firms, or markets (Gopalan et al., 2021; Bassier, 2022; Demir, 2024; Giupponi and Machin, 2026). Other work studies how information and search assistance affect mobility across occupations and employers (Jäger et al., 2024; Altmann et al., 2023; Belot et al., 2019, 2025, 2026; Klaeui, 2026). Our contribution is to estimate an occupation-level wage response and to provide a framework for separating the response in the treated occupation from the displacement of applications elsewhere on the platform.

Section 2 describes the institutional setting and data. Section 3 develops the conceptual framework. Sections 4 and 5 present the reduced-form evidence on applications and hires. Section 6 estimates the structural model and decomposes the reduced-form effects. Section 7 concludes.

2 Institutional context and data

2.1 The online platform

Extracadabra is an online platform in the French labour market that specialises in hospitality jobs in and around Paris. Its clients are mainly restaurants, bars, hotels, and event organisers. The platform advertises three types of arrangement: standard employment contracts of fixed or indefinite duration, temporary-employment contracts, and self-employed contractor arrangements in which workers supply services through their own businesses. We focus on the last category, which is the most common on the platform and on which the richest data are available. These assignments typically last from a few hours to a few days and must be filled at short notice. In the French hospitality industry, such occasional shifts are known as *extras*.

Extracadabra differs from a conventional job board because workers do not search the full set of posted vacancies. An employer first provides the characteristics of the position to the

platform (Appendix A.1). A matching algorithm then identifies a set of suitable workers and sends each of them a notification through the mobile application. Notified workers can inspect the job details and either ignore the notification, decline the opportunity, or apply (Appendix A.2). The employer is notified of each application, reviews the applicant profiles, and selects one or more workers or leaves the position unfilled. Once a worker is selected, the platform charges the employer the wage bill and its commission and holds the payment in escrow.

After the job is completed, the employer can rate the worker on a scale from 1 to 5. Ratings are usually either missing or equal to 4 or 5. A worker rated 4 or above becomes an approved worker for that employer–occupation pair and receives priority in the matching algorithm for the employer’s subsequent vacancies in the same occupation.

The platform also plays an unusually direct role in wage setting. It specifies a default hourly wage for each occupation. Employers can depart from this default by asking customer service to create a client-specific wage grid. Since 19 December 2023, they have also been able to boost the wage when creating a position. Between April 2023 and May 2024, 94.6% of posted wages equalled the occupation-wide default, 4.5% equalled a client-specific default, and 0.9% used the wage-boost option. Our empirical design exploits a change in one of these occupation-specific defaults.

2.2 Data sources

The platform provides job-level, worker-level, and worker–job-level data.

Job-level data

For each job, we observe the occupation, workplace postcode, complete schedule (including unpaid breaks), number of workers required, and the date and time at which the position was created. We also observe the employer’s identity, location, sector, and, for bars and restaurants, a detailed establishment category such as “Italian restaurant”.

Worker-level data

Workers must register before applying. Their profiles contain age, gender, postcode of residence, declared skills, and previous work experience. The platform updates these profiles as workers complete assignments, so at each date we observe a worker’s prior assignments and the associated employer ratings. We also observe every time a worker connects to the mobile application.

Pair-level data

A distinctive feature of the data is that they cover each stage of the matching process. For every job, we observe the workers notified by the platform, including those who do not apply. We then observe each application and its timestamp, whether and when the employer ignored, rejected, or selected the applicant, and the workers who ultimately performed the job. For completed jobs, we observe hours worked and the wage bill transferred through the platform.

2.3 Sample and summary statistics

The analysis sample contains jobs posted between 5 April 2023 and 4 May 2024 for waiting staff and three neighbouring occupations: bar staff, runners, and receptionists. We exclude the small minority of jobs that platform staff intermediated manually.

The sample contains 26,297 job adverts posted by 1,475 firms at 2,200 venues, where a venue is an establishment–location pair. At least one worker was hired for 47.1% of these jobs. The worker sample comprises the 5,632 workers who submitted at least one application to a sample job; together, they made 320,833 applications.

Table 1 reports job characteristics by occupation. Jobs last an average of 6 hours, and only 20% span more than one day. They include an average of 30 minutes of unpaid breaks and are posted 3.5 days before they begin. Employers request an average of 1.3 workers per job. Almost all jobs are in the Paris region, and 66–88% are within the city of Paris, depending on the occupation. The platform notifies an average of 774–1,335 workers per job, of whom 8.2–19.4 apply. Employers hire an average of 0.55 workers per job, and approximately half of all jobs result in at least one hire.

Table 1: Summary statistics for jobs

	Waiting staff	Bar staff	Receptionist	Runner	Non-waiting staff
Number notified	1,335.09 (496.71)	898.50 (297.72)	1,159.87 (471.15)	773.99 (197.59)	885.07 (314.65)
Number of applications	13.77 (12.55)	10.61 (9.63)	19.38 (13.67)	8.15 (8.31)	10.74 (10.27)
Number of hires	0.55 (0.71)	0.61 (0.78)	0.59 (0.73)	0.49 (0.66)	0.57 (0.74)
At least one hire	0.46 (0.50)	0.50 (0.50)	0.48 (0.50)	0.43 (0.50)	0.48 (0.50)
Number of workers requested	1.31 (1.27)	1.27 (1.30)	1.35 (1.17)	1.23 (1.30)	1.27 (1.28)
Work hours	5.75 (1.94)	5.86 (1.82)	6.08 (2.03)	6.91 (1.85)	6.24 (1.92)
Break hours	0.30 (0.61)	0.33 (0.55)	0.39 (0.49)	0.40 (0.55)	0.36 (0.55)
Time until shift start (hours)	88.91 (141.89)	58.26 (97.83)	112.43 (151.39)	111.84 (177.76)	82.28 (138.49)
Located in Paris	0.66 (0.47)	0.69 (0.46)	0.88 (0.32)	0.71 (0.45)	0.72 (0.45)
Located in IDF	0.96 (0.21)	0.91 (0.28)	0.99 (0.12)	0.85 (0.35)	0.90 (0.30)
Morning	0.14 (0.35)	0.10 (0.30)	0.23 (0.42)	0.07 (0.26)	0.11 (0.31)
Lunch	0.44 (0.50)	0.43 (0.49)	0.48 (0.50)	0.26 (0.44)	0.37 (0.48)
Afternoon	0.27 (0.44)	0.31 (0.46)	0.39 (0.49)	0.27 (0.44)	0.31 (0.46)
Evening	0.34 (0.47)	0.30 (0.46)	0.32 (0.47)	0.24 (0.43)	0.28 (0.45)
Night	0.25 (0.43)	0.33 (0.47)	0.24 (0.43)	0.57 (0.49)	0.40 (0.49)
Number of working days					
1	0.81	0.79	0.74	0.76	0.77
2	0.09	0.10	0.11	0.12	0.11
3	0.03	0.04	0.06	0.04	0.04
4	0.02	0.02	0.02	0.02	0.02
5	0.03	0.03	0.05	0.04	0.04
>5	0.02	0.02	0.03	0.02	0.02
Type of shift					
Week	0.64	0.55	0.60	0.56	0.56
Weekend	0.06	0.09	0.12	0.08	0.09
Week & Weekend	0.30	0.36	0.29	0.36	0.35
Share of sample	0.48	0.29	0.06	0.17	0.52
N	12,677	7,530	1,491	4,599	13,620

Notes: The table reports summary statistics for jobs posted between 5 April 2023 and 4 May 2024 in the waiting-staff, bar-staff, receptionist, and runner occupations. Standard deviations are reported in parentheses below means. No additional statistic is reported below proportions.


Figure 1 shows the number of jobs created each month. Postings are strongly seasonal: demand is highest in June and September and is approximately twice as high as in July and August. Workers do not observe the employer's name when applying, which helps prevent employers and workers from bypassing the platform. Instead, they see a detailed category describing the type of establishment. Figure 2 displays the distribution of establishment categories. The largest categories are café-restaurants, Italian restaurants, fast-food establishments, three-star hotels, and catering services.

Table 2 reports applicant characteristics in Panel A and employer characteristics in Panel B. Because workers can apply in several occupations, the occupation-specific columns are not mutually exclusive. Waiting-staff applicants are 28 years old on average, and 54% are women. Almost all (99%) live in Île-de-France, and 74% live within the city of Paris. Their homes are an average of 17 kilometres from the job: 29% live in the same *département* as the workplace and 1% in the same postcode. Applicants are highly active on the platform. Before receiving a job notification, 56% had connected within the previous hour and 86% within the previous 12 hours. Their average rating is 4.5, and 0.8% have no previous rating.

Applicant characteristics are broadly similar across occupations. Gender is a notable exception: receptionist applicants are more likely to be women, bar-staff and runner applicants are more likely to be men, and waiting-staff applicants lie between these groups.


Panel B describes the 1,475 firms in the sample. Because firms can post jobs in several occupations, they may also appear in several columns. Firms posting waiting-staff jobs operate an average of 1.5 venues, compared with approximately 1.2 among firms posting in the other three occupations. Conditional on posting in an occupation, firms create an average of 12 adverts during the sample period.

Figure 1: Monthly job postings



Notes: The figure reports monthly job postings for waiting staff, bar staff, receptionists, and runners from 5 April 2023 to 4 May 2024.

Figure 2: Job postings by establishment category



Notes: The figure reports the distribution of job postings across establishment categories for waiting staff, bar staff, receptionists, and runners from 5 April 2023 to 4 May 2024. Categories with fewer than 50 postings are grouped as “Other”.

Table 2: Summary statistics of workers and firms

	Waiting staff	Bar staff	Receptionist	Runner	Non-waiting staff
Panel A: applicant characteristics					
Age	27.96 (6.52)	28.55 (6.42)	26.98 (5.41)	29.06 (6.89)	28.37 (6.40)
Female	0.54 (0.50)	0.36 (0.48)	0.71 (0.45)	0.39 (0.49)	0.44 (0.50)
Avg. rating	4.53 (0.58)	4.53 (0.56)	4.55 (0.60)	4.50 (0.63)	4.53 (0.59)
Located in IDF	0.99 (0.10)	0.97 (0.16)	1.00 (0.07)	0.92 (0.27)	0.97 (0.18)
Located in Paris	0.74 (0.44)	0.78 (0.42)	0.90 (0.30)	0.80 (0.40)	0.81 (0.39)
Crow distance	17.31 (44.95)	17.85 (53.39)	15.39 (47.70)	17.24 (51.94)	17.21 (51.95)
Same department	0.29 (0.45)	0.34 (0.48)	0.33 (0.47)	0.33 (0.47)	0.34 (0.47)
Same postcode	0.01 (0.12)	0.02 (0.14)	0.02 (0.13)	0.02 (0.13)	0.02 (0.13)
Time since last connection					
Connected in the last hour	0.56	0.58	0.51	0.50	0.54
Connected in the 6 hours	0.21	0.22	0.22	0.23	0.22
Connected in the 12 hours	0.09	0.08	0.09	0.09	0.09
Connected in the last day	0.08	0.07	0.10	0.09	0.08
Connected in last week	0.05	0.04	0.07	0.07	0.06
No recent connection	0.01	0.01	0.01	0.01	0.01
N distinct individuals	4,007	2,307	2,519	1,716	4,539
N	174,577	79,894	28,895	37,467	146,256
Panel B: firm characteristics					
Number of venues	1.52 (2.44)	1.21 (1.18)	1.22 (1.04)	1.19 (0.89)	1.20 (1.05)
Number of postings	11.51 (28.54)	13.94 (30.13)	10.57 (19.45)	10.22 (21.57)	12.04 (25.85)
N distinct establishments	1,101	540	141	450	848
N	12,677	7,530	1,491	4,599	13,620

Notes: The table reports applicant and employer characteristics for waiting-staff, bar-staff, receptionist, and runner positions advertised between 5 April 2023 and 4 May 2024. Standard deviations are reported in parentheses below means. No additional statistic is reported below proportions.


Figure 3 plots monthly median and mean hourly wages in the four occupations. Because most employers use the platform default, the median generally equals the default wage.

The default remained unchanged throughout the sample period for receptionists, runners, and bar staff. On 5 October 2023, however, the waiting-staff default rose from €16 to €17, an increase of 6.25%. The platform notified client firms of the change.

Mean wages are noisier because employers can use client-specific wage grids or boost individual postings. Boosting is particularly visible for bar-staff and receptionist jobs in December 2023. Nevertheless, both panels show a clear increase in waiting-staff wages relative to wages in the three comparison occupations.

We use this relative wage change to estimate effects on the quantity and quality of applications and on hiring.

Figure 3: Hourly wages by month and occupation



Notes: The figure reports median and mean advertised hourly wages for waiting-staff, bar-staff, receptionist, and runner jobs between 5 April 2023 and 4 May 2024. Values are calculated from the job-level data.

3 Conceptual framework

We develop a model of application and hiring decisions on the platform. The model serves two purposes. First, it clarifies why a higher wage need not increase applications from every type of worker. A worker applies to a job not only because it pays well, but also because the worker has a reasonable chance of obtaining it. A wage increase can attract additional high-productivity applicants, lower the hiring prospects of other applicants, and discourage some of them from applying. Second, the model shows how a wage change in one occupation affects applications to other occupations through the value of continuing to search.

We abstract from several institutional details. The framework contains two occupations: occupation A , which represents waiting-staff jobs and whose wage changes in the empirical analysis, and occupation B , which represents the pooled comparison occupations. In Section 6, we estimate the model's parameters and use them to decompose the reduced-form estimates into their equilibrium components. Appendix C provides more detail about the conceptual framework, as well as the proofs.

3.1 Environment

Time is discrete. Workers are risk neutral and discount future utility by $\tilde{\beta}$. A non-employed worker has type

$$q = (q_A, q_B) \in [0, 1]^2,$$

where q_j is the worker's productivity rank in occupation j . Higher values indicate higher productivity. The joint distribution $F(q_A, q_B)$ allows workers who are productive in one occupation to be systematically more or less productive in the other.

In each period, a non-employed worker receives at most one notification. A notification from occupation $j \in \{A, B\}$ arrives at a type-specific rate $\alpha_j(q)$, which is observed from pre-reform notification exposure. Upon receiving it, the worker observes a match-specific payoff shock ε_j , distributed logistically with scale s_j , and decides whether to apply. Applying costs $\tilde{c}_j$. Employment in occupation j pays $w_j + \varepsilon_j$ per period and ends with probability δ . After their employment spell, the worker returns to search.

Each vacancy in occupation j notifies M_j workers. Employers value both observed productivity and an idiosyncratic employer–applicant match component. Specifically, employer j 's payoff from applicant i is $q_{ij} + v_{ij}$, where v_{ij} follows a type-I extreme-value distribution with scale σ_j . The employer selects the applicant with the highest payoff, and

the selected applicant is successfully matched with probability κ_j . Thus σ_j governs how strongly employers sort applicants by observed productivity: selection approaches deterministic ranking when σ_j is small and random selection when it is large. We derive the resulting selection probabilities below.

3.2 Application decisions

Let $U(q; w)$ be the value of remaining non-employed before the worker receives a notification, where $w = (w_A, w_B)$, and define the flow-equivalent continuation value

$$u(q; w) = (1 - \tilde{\beta})U(q; w).$$

The flow value of remaining available for work is $b(q) = b + b_1 \bar{q}$, where $\bar{q} = (q_A + q_B)/2$. This specification allows the value of continued search to vary with a worker's average productivity rank. Let $p_j(q_j; w)$ denote the equilibrium probability that a worker is hired conditional on applying to occupation j . After rescaling the application cost by the expected duration of employment, the worker applies if and only if

$$p_j(q_j; w) [w_j + \varepsilon_j - u(q; w)] \geq c_j. \quad (1)$$

Conditional on a notification from occupation j , a worker's application probability is

$$\lambda_j(q; w) = \frac{1}{1 + \exp \{ -[w_j - u(q; w) - c_j / p_j(q_j; w)] / s_j \}}. \quad (2)$$

Under equation (2), the application behaviour is different from a standard reservation-wage model. With a fixed hiring probability, the worker compares the wage with a reservation value. Here that reservation value is itself higher when the job is harder to obtain:

$$\text{effective reservation value} = u(q; w) + \frac{c_j}{p_j(q_j; w)}.$$

A low-productivity worker may therefore decline a high-wage opportunity if the expected cost of securing the job is sufficiently large.

The continuation value links the two occupations. With logistic shocks it can be written

$$u(q; w) = b(q) - \sum_{k \in \{A, B\}} \alpha_k(q) \beta s_k p_k(q_k; w) \log[1 - \lambda_k(q; w)], \quad \beta = \frac{\tilde{\beta}}{1 - \tilde{\beta}(1 - \delta)}. \quad (3)$$

Better opportunities in A raise the value of waiting for an A notification. This makes a current opportunity in B less attractive, even though w_B has not changed.

3.3 Equilibrium competition

Application decisions determine the applicant pool and hence the probability of being hired. The quantitative implementation solves the model on the empirical cells c of the joint (q_A, q_B) distribution. Let ω_c be the population share of cell c , and let $v_j(c) \propto \omega_c \alpha_j(c)$ be the distribution of workers notified by occupation j . Because the match component v_{ij} is type-I extreme-value, the probability that an employer selects an applicant with productivity q_j , conditional on a realised applicant pool, is the applicant's index $\exp(q_j/\sigma_j)$ divided by the sum of these indices over all applicants. We approximate the random sum of competitors' indices by its expectation. The resulting hiring probability is

$$p_j(q_j; w) = \kappa_j \frac{\exp(q_j/\sigma_j)}{\exp(q_j/\sigma_j) + D_j(w)}, \quad (4)$$

$$D_j(w) = M_j \sum_c v_j(c) \exp(q_j(c)/\sigma_j) \lambda_j(c; w). \quad (5)$$

The index $D_j(w)$ summarizes effective competition from other applicants. An equilibrium consists of application probabilities, hiring probabilities, and continuation values satisfying equations (2)–(5).¹ The wage response of applications combines three forces.

Proposition 1 (Wages, competition, and application decisions). *For a worker of type q ,*


$$\frac{d\lambda_j(q; w)}{dw_k} = \frac{\lambda_j(q; w)[1 - \lambda_j(q; w)]}{s_j} \left[\mathbf{1}\{j = k\} - \frac{du(q; w)}{dw_k} + \frac{c_j}{p_j(q_j; w)^2} \frac{dp_j(q_j; w)}{dw_k} \right]. \quad (6)$$

Consequently, an increase in w_j need not increase applications to occupation j for every worker type.

For the occupation whose wage rises, the first term in brackets is the direct wage effect. The second term is an option-value effect: a higher wage raises the value of remaining available for future opportunities and therefore makes the current job less attractive. The third term is the competition effect. If the wage attracts additional applicants, D_j rises and p_j falls. The resulting increase in c_j/p_j is especially important for workers with low hiring probabilities. For these workers, the competition and option-value effects can outweigh the direct wage effect.

¹The mean-field representation is an approximation. The exact Poisson-logit hiring operator and the full equilibrium definition are given in Appendix C.

Figure 4: Wages, hiring prospects, and application decisions



Notes: The figure solves the model under an illustrative calibration. Panel (a) reports the average probability of applying to occupation A , conditional on an A notification, at two wages. Panel (b) reports the same probability as a function of w_A for three productivity ranks. Continuation values and hiring probabilities adjust in equilibrium along each curve.

Figure 4 illustrates this possibility. Panel (a) compares application profiles at two wages. The profiles cross because the wage increase raises applications among high-productivity workers but lowers them for some lower-productivity workers. Panel (b) follows three worker types over a wider wage range. Applications rise monotonically for the high-productivity type. For the two lower-productivity types, applications first rise with the wage and then fall once the induced increase in competition dominates the direct wage incentive. The numerical values are chosen only to display the mechanism and are not estimates.

The same expression clarifies spillovers. When w_A rises, applications to B have no direct wage term:

$$\frac{d\lambda_B}{dw_A} = \frac{\lambda_B(1 - \lambda_B)}{s_B} \left[-\frac{du}{dw_A} + \frac{c_B}{p_B^2} \frac{dp_B}{dw_A} \right].$$

The increase in the continuation value tends to reduce applications to B . As those applications fall, competition for B jobs weakens and p_B rises, which mitigates the initial decline. The size and composition of the spillover depend on the joint distribution of q_A and q_B : it is strongest when workers who are marginal applicants to B also have good prospects in A .

3.4 Observable outcomes

The model maps application decisions into the three classes of outcomes studied empirically. Expected applications per vacancy are

$$\mathcal{A}_j(w) = M_j \sum_c v_j(c) \lambda_j(c; w). \quad (7)$$

The productivity distribution among applicants is the notified type distribution reweighted by $\lambda_j(c; w)$. In particular, mean applicant quality is

$$\bar{Q}_j(w) = \frac{\sum_c v_j(c) q_j(c) \lambda_j(c; w)}{\sum_c v_j(c) \lambda_j(c; w)}. \quad (8)$$

The appendix gives the corresponding expressions for other functionals of the applicant pool, including its minimum and maximum quality.


Finally, Poisson application arrivals imply that the job-filling rate is

$$H_j(w) = \kappa_j [1 - \exp\{-\mathcal{A}_j(w)\}]. \quad (9)$$

Applications therefore raise filling mechanically, but the marginal effect becomes small when vacancies already receive many applications. Applicant quality has no similarly mechanical response: it rises when the wage draws disproportionately more high-productivity workers or discourages low-productivity workers, and falls under the opposite compositional change.

Figure 5 illustrates these mappings. In the example, a higher w_A strongly raises applications and average applicant quality in A , while slightly reducing applications to B . These equilibrium links matter for the difference-in-differences design. The observed change for occupation A includes the direct wage incentive and the endogenous responses of competition and continuation values. Moreover, occupation B need not be unaffected: the improved opportunity in A changes workers' outside option and can alter applications, applicant quality, and filling in B . The structural analysis uses the estimated model to separate these components.

Figure 5: From wages to applications and applicant quality



Notes: The figure uses the same illustrative calibration as Figure 4. Occupation B's wage is held fixed. Applicant quality is the mean occupation-specific productivity rank among applicants.

4 The impact of wages on applications

This section estimates the effect of the waiting-staff wage increase on the quantity and quality of applications, using the other three occupations as the comparison group.

4.1 Empirical strategy

We compare outcomes before and after 5 October 2023 for waiting-staff jobs and jobs in the three comparison occupations. Because employers can depart from the platform default, the estimates are intention-to-treat effects.²

For the number of applications A_j received by job j , we estimate the following Poisson model:

$$\mathbb{E}[A_j \mid X_j] = \exp\left(\beta Z_{o(j)} \mathbb{1}\{t(j) > t_0\} + X_j' \alpha + \lambda_{o(j)} + \lambda_{\text{week}(j)} + \phi_j + \ln M_j\right). \quad (10)$$

For an applicant-quality outcome Y_j , we estimate the corresponding linear model:

$$Y_j = \beta Z_{o(j)} \mathbb{1}\{t(j) > t_0\} + X_j' \alpha + \lambda_{o(j)} + \lambda_{\text{week}(j)} + \phi_j + \varepsilon_j. \quad (11)$$

Here $o(j)$ and $t(j)$ denote the occupation and creation date of job j , $Z_{o(j)}$ equals one for waiting-staff jobs, and t_0 is the date of the wage change. The vector X_j includes indicators for morning, lunch, afternoon, evening, and night shifts; contracted work and break hours; the day of the month, day of the week, and hour at which the advert was created; deciles of the time between posting and the start of the job; indicators for the number of working days spanned by the job; and indicators for weekday-only, weekend-only, and mixed jobs. In equation (10), $\ln M_j$ is an offset for the number of workers notified about job j .

All specifications include occupation and week fixed effects. The additional fixed effects in ϕ_j vary across columns. Our preferred specification includes venue fixed effects, where a venue is an establishment–postcode pair. The coefficient β captures the intention-to-treat effect of the default-wage increase. For quality, we consider the mean, minimum, and maximum score among applicants to each job.

²In the analysis sample, 95% of adverts post the occupation-wide default wage, 2% post a lower wage, and 3% post a higher wage.

We also estimate event-study versions of equations (10) and (11):

$$\mathbb{E}[A_j | X_j] = \exp \left(\sum_{k=-6, k \neq -1}^6 \beta_k Z_{o(j)} \mathbb{1}\{m(j) = k\} + X_j' \alpha + \lambda_{o(j)} + \lambda_{\text{week}(j)} + \phi_j + \ln M_j \right) \quad (12)$$

$$Y_j = \sum_{k=-6, k \neq -1}^6 \beta_k Z_{o(j)} \mathbb{1}\{m(j) = k\} + X_j' \alpha + \lambda_{o(j)} + \lambda_{\text{week}(j)} + \phi_j + \varepsilon_j \quad (13)$$

where $m(j)$ indexes event months relative to t_0 ; the month from 5 October to 4 November 2023 is zero, the next month is one, and so forth. The omitted category is the month immediately preceding the change. Defining months around the event date ensures that the default wage is constant within each event month.

Identification requires that, absent the wage change, outcomes in the treated and comparison occupations would have followed parallel trends. Occupation-specific trends or seasonal patterns could violate this condition. Detailed job controls and venue fixed effects reduce concerns about changes in the composition of postings. Moreover, many of the same firms post jobs, and many of the same workers apply, across the four occupations. Common demand- or supply-side shocks are therefore likely to affect more than one occupation.

This overlap also creates interference: additional applications to waiting-staff jobs may come at the expense of the comparison occupations. The reduced-form coefficient consequently identifies a relative effect, not necessarily the equilibrium response in the treated occupation alone. Section 6 uses the model to separate these objects.

4.2 Defining quality

We construct quality from employers' revealed preferences over applicants. Specifically, we use a class-weighted logistic model to predict whether worker i is hired for job j , conditional on applying. Because the estimation reweights hires and non-hires to address class imbalance, the fitted value is best interpreted as a hiring score rather than as a calibrated probability. Quality is worker-job specific: the same worker can have different scores for different jobs.


$$S_{ij} = \frac{1}{1 + \exp(-z_{ij})}, \quad z_{ij} = \beta_0 + X_{ij}' \beta + \gamma_{c(j), e(j)}, \quad (14)$$

where S_{ij} is the quality score, X_{ij} contains worker, job, and worker–job characteristics, and $\gamma_{c(j),e(j)}$ denotes establishment fixed effects indexed by establishment category and employer. The predictors include whether the worker has previously worked for and been approved by the employer, distance between the worker’s residence and the workplace, the worker’s average rating, experience in the job’s occupation and establishment category, gender, age, and time since the worker last connected to the platform. We estimate the model using service-sector applications between 1 January 2023 and 30 December 2025. Appendix B describes the sample, predictors, and estimation in detail.

We draw 20% of jobs at random to form the training sample and calculate scores for applications to the remaining 80%. The quality regressions in equations (11) and (13) exclude all jobs in the training sample, preventing a job’s realised hiring outcome from being used both to estimate and to evaluate its applicants’ quality scores.

In the test sample, the area under the receiver operating characteristic curve (AUC) is 0.76. At the selected threshold, sensitivity—the share of hires classified as hires—is 0.61. Figure 6 plots the score distribution separately for hired and non-hired applicants. Although the distributions overlap, hired applicants receive substantially higher scores.

Figure 6: Applicant quality scores by hiring outcome



Notes: The figure reports the density of applicant quality scores by hiring outcome in the test sample. The sample contains applications to service-sector jobs posted between 1 January 2023 and 30 December 2025, excluding jobs used to train the quality model.

4.3 Results

Table 3 reports estimates of equation (10). Column (1) includes occupation and week fixed effects; column (2) adds establishment-category and postcode fixed effects; column (3) adds

their interaction; and column (4) instead includes venue fixed effects. Across specifications, the wage increase raised applications by 22–30% relative to the comparison occupations.³ In our preferred specification, $\hat{\beta} = 0.2104$, corresponding to a 23.4% increase. The log elasticity is $0.2104 / \log(17/16) = 3.5$.

Table 3: Effect of the wage increase on applications

	Number of applications			
	(1)	(2)	(3)	(4)
Post $\times$ Treated	0.2662*** (0.0204)	0.1950*** (0.0200)	0.2066*** (0.0209)	0.2104*** (0.0221)
Observations (jobs)	26,297	26,112	25,921	25,461
Share treated	0.482	0.48	0.478	0.474
Dependent-variable mean	12.2	12.2	12.3	12.2
Number of Venues	2,200	2,021	1,832	1,379
Controls	✓	✓	✓	✓
Occupation FE	✓	✓	✓	✓
Week FE	✓	✓	✓	✓
Category FE		✓	✓	
Postcode FE		✓	✓	
Category $\times$ Postcode FE			✓	
Venue FE				✓


Notes: The sample runs from 5 April 2023 to 4 May 2024. All specifications are Poisson models with the number of notified workers as an offset. Heteroskedasticity-robust standard errors are in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure 7 presents the event-study estimates from the preferred specification of equation (12). The pre-treatment coefficients are small and statistically insignificant, providing no evidence of differential application trends before the wage change. Applications rise gradually after the change, reaching approximately 20% above the relative pre-treatment level after two months. The effect persists through the end of the six-month post-treatment period.

³Percentage changes are calculated as $100[\exp(\hat{\beta}) - 1]$.

Figure 7: Event study: number of applications



Notes: The figure reports the estimated β_k coefficients from equation (12). The model includes occupation, week, and venue fixed effects and uses the number of notified workers as an offset. The sample runs from 5 April 2023 to 4 May 2024. Bars show 95% confidence intervals.

Table 4 reports effects on the mean, minimum, and maximum quality scores in the applicant pool. Mean quality increases by 0.007–0.013 across specifications. In the preferred specification, the estimate is 0.0095, or 2.0% of the dependent-variable mean, implying an elasticity of approximately 0.3. The minimum score rises by 0.0307, or 10.6% of its mean, implying an elasticity of approximately 1.7. By contrast, the maximum-score estimate is 0.0065, or 0.9% of its mean, and is statistically insignificant. The improvement in applicant quality is therefore concentrated at the lower end of the distribution.

Table 4: Effect of the wage increase on applicant quality

Applicant	Mean quality				Minimum quality				Maximum quality			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Post $\times$ Treated	0.0065** (0.0031)	0.0128*** (0.0030)	0.0107*** (0.0032)	0.0095*** (0.0034)	0.0292*** (0.0036)	0.0349*** (0.0037)	0.0329*** (0.0039)	0.0307*** (0.0043)	0.0042 (0.0052)	0.0095* (0.0051)	0.0060 (0.0053)	0.0065 (0.0057)
Observations (jobs)	19,503	19,388	19,207	18,773	19,503	19,388	19,207	18,773	19,503	19,388	19,207	18,773
Share treated	0.491	0.49	0.488	0.483	0.491	0.49	0.488	0.483	0.491	0.49	0.488	0.483
Dependent-variable mean	0.484	0.484	0.485	0.485	0.289	0.288	0.289	0.289	0.726	0.726	0.727	0.728
Number of Venues	1,915	1,800	1,619	1,185	1,915	1,800	1,619	1,185	1,915	1,800	1,619	1,185
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Occupation FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Week FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Category FE		✓	✓			✓	✓			✓	✓	
Postcode FE		✓	✓			✓	✓			✓	✓	
Category $\times$ Postcode FE			✓				✓				✓	
Venue FE				✓				✓				✓


Notes: The sample runs from 5 April 2023 to 4 May 2024. All specifications are estimated by OLS. Heteroskedasticity-robust standard errors are in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.


Figure 8 presents the event-study counterparts of columns (4), (8), and (12) of Table 4.

The increase in mean quality is short-lived and statistically significant only in the first two months after the event month. Minimum quality rises more sharply and remains significantly higher throughout the post-treatment period. Maximum quality shows no clear or persistent response.


Figure 8: Event study: applicant quality



(a) Mean applicant quality



(b) Minimum applicant quality



(c) Maximum applicant quality

Notes: Panels 8a–8c report the estimated β_k coefficients from equation (13) for mean, minimum, and maximum applicant quality. All specifications include occupation, week, and venue fixed effects. The sample runs from 5 April 2023 to 4 May 2024. Bars show 95% confidence intervals.

5 The impact of wages on hires

We apply the same difference-in-differences and event-study designs to the number of workers hired for each job. Because the default wage is not binding, these estimates are again intention-to-treat effects.

Identification requires parallel trends in hiring absent the wage change. The threats are similar to those discussed in Section 4, and we use the same job controls and fixed effects. Hiring differs from applying, however, because it reflects the joint decisions of workers and employers. Employers may respond to the wage increase by changing their selection standards or recruitment practices. The estimated effect on hires therefore combines the application response with any wage-induced change in employers' hiring decisions.

Table 5 reports the effect on the number of hires. Across specifications, the wage increase raises hires by 11–16% relative to the comparison occupations. In the preferred specification, $\hat{\beta} = 0.1087$, corresponding to an 11.5% increase. The implied log elasticity is 1.8. Figure 9 presents the corresponding event study. The post-treatment coefficients are generally positive but individually imprecise, so none is statistically distinguishable from zero. The pooled post-treatment estimate in Table 5, by contrast, is positive and statistically significant.


Table 5: Effect of the wage increase on hires

	Number of hires			
	(1)	(2)	(3)	(4)
Post $\times$ Treated	0.1028*** (0.0338)	0.1086*** (0.0342)	0.1481*** (0.0354)	0.1087*** (0.0375)
Observations (jobs)	26,297	25,938	25,509	24,629
Share treated	0.482	0.478	0.476	0.471
Dependent-variable mean	0.6	0.6	0.6	0.6
Number of Venues	2,200	1,952	1,694	1,137
Controls	✓	✓	✓	✓
Occupation FE	✓	✓	✓	✓
Week FE	✓	✓	✓	✓
Category FE		✓	✓	
Postcode FE		✓	✓	
Category $\times$ Postcode FE			✓	

Notes: The sample runs from 5 April 2023 to 4 May 2024. All specifications are Poisson models with the number of notified workers as an offset. Heteroskedasticity-robust standard errors are in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure 9: Event study: number of hires



Notes: The figure reports the estimated β_k coefficients from equation (12) for the number of hires. The model includes occupation, week, and venue fixed effects and uses the number of notified workers as an offset. The sample runs from 5 April 2023 to 4 May 2024. Bars show 95% confidence intervals.

6 Structural decomposition of the reduced-form estimators

We estimate the model developed in Section 3. Our main purpose is to decompose the difference-in-differences estimates into the equilibrium effect on the occupation in which the wage increased and the displacement effect on the comparison occupations. A secondary benefit is that the richness of the platform data identifies structural parameters that are typically difficult to recover, including application costs and the dispersion in workers' and employers' evaluations of a match. Details about the estimation and the decomposition are available in Appendix D.

6.1 Parameters and moments

The model contains three groups of parameters. The first group is observed or calibrated directly:

$$\theta_0 = (M_A, M_B, \{\alpha_A(c), \alpha_B(c)\}_c, \delta, \tilde{\beta}, \{\omega_c, q_A(c), q_B(c)\}_c).$$

The numbers of notified workers per vacancy M_j , cell-specific notification rates $\alpha_j(c)$, cell weights, representative cell ranks, and mean match duration are observed on the platform. The discount factor is fixed from an annual discount rate. Earlier versions of the quantitative model used a Gaussian copula for the joint distribution of ranks; the current baseline instead uses the empirical joint cells directly.

The second group contains occupation-specific matching efficiencies (κ_A, κ_B) . Let $\mathcal{A}_j$ denote applications per vacancy and z_j the fraction of vacancies that remain unfilled. Because application arrivals are Poisson,

$$\kappa_j = \frac{1 - z_j}{1 - \exp(-\mathcal{A}_j)}. \quad (15)$$

We recover κ_j from pre-reform applications and filling rates.

The remaining behavioural parameters are estimated jointly:

$$\theta = (b, b_1, c_A, c_B, s_A, s_B, \sigma_A, \sigma_B)'. \quad (16)$$

The intercept b governs the attractiveness of remaining available for future work, and b_1 allows this value to vary with the worker's average productivity rank. The c_j 's are application costs, the s_j 's govern dispersion in workers' application payoffs, and the σ_j 's govern dispersion in employers' evaluations of applicants.

We construct pre-reform moments by crossing bins of q_A and q_B . For every joint worker cell c , we measure the probability of applying conditional on receiving a notification:

$$\hat{\lambda}_j(c) = \frac{\text{applications from cell } c \text{ to } j}{\text{notifications received by cell } c \text{ from } j}.$$

We also measure hiring probabilities conditional on applying. Consistent with the model, p_A is estimated by q_A bin and p_B by q_B bin:

$$\hat{p}_j(x) = \frac{\text{hires from applications in rank bin } x}{\text{applications in rank bin } x}.$$

The pre-reform moments additionally include applications per vacancy in each occupation. Appendix D describes the sample, weights, standard errors, and construction of the representative comparison occupation.

Finally, we add the empirical difference-in-differences estimates for applications and average applicant quality. Let w_A^0 be the pre-reform wage and $w_A^1 = w_A^0 + \Delta w_A$ the post-reform wage. Solving the model at both wages gives

$$\tau_A^M(\theta) = [h(\mathcal{A}_A^1) - h(\mathcal{A}_A^0)] - [h(\mathcal{A}_B^1) - h(\mathcal{A}_B^0)], \quad (17)$$

$$\tau_Q^M(\theta) = [\bar{Q}_A^1 - \bar{Q}_A^0] - [\bar{Q}_B^1 - \bar{Q}_B^0], \quad (18)$$

where h applies the same transformation as the empirical specification. Because the application regression is Poisson, its coefficient is a change in the log application rate, so its structural counterpart uses $h(x) = \log x$. Applicant quality is expressed in occupation-specific percentile-rank units in both the data and the model.

6.2 Identification

Identification follows the sequence suggested by the economic structure. Duration and notification moments identify $(\delta, \tilde{\beta}, \alpha_j(c), M_j)$, filling and application rates identify κ_j through equation (15), and the empirical cell weights identify the joint distribution of productivity ranks used in the equilibrium computation.

The slope of hiring probabilities across productivity ranks identifies employer selectivity. Rearranging the mean-field hiring equation yields

$$\log \left[\frac{\kappa_j}{p_j(q_j)} - 1 \right] = \log D_j - \frac{q_j}{\sigma_j}. \quad (19)$$

Thus the slope of transformed hiring probabilities in q_j is $-1/\sigma_j$. A flatter hiring profile implies more employer-side noise and a larger σ_j .

Conditional application probabilities identify the worker-side parameters. Define

$$G_j(q) = -\alpha_j(q)\beta p_j(q_j) \log[1 - \lambda_j(q)].$$

Combining the application rule with the continuation-value equation gives

$$w_j = b + b_1 \bar{q} + \frac{c_j}{p_j(q_j)} + s_j \text{logit}[\lambda_j(q)] + s_A G_A(q) + s_B G_B(q). \quad (20)$$

Across occupations and productivity cells, equation (20) is linear in $(b, b_1, c_A, c_B, s_A, s_B)$. These parameters are point identified when the resulting regressor matrix has rank six. Intuitively, variation in hiring probabilities identifies application costs, variation in application rates identifies the worker-side shock scales, and behaviour across both occupations and quality cells identifies the level and type gradient of the continuation value.

The pre-reform conditional moments are sufficient for point identification under these rank conditions. The two difference-in-differences moments add economically important overidentifying information. They discipline the parameter combinations that determine continuation-value and competition responses to wages and require the model's comparative statics to align with the reduced-form evidence. Appendix D states the formal identification result.

6.3 GMM estimator

Let $\hat{m}$ collect the pre-reform cell and aggregate moments and the two difference-in-differences estimates, and let $m(\theta)$ denote their model counterparts. We standardise each discrepancy by its estimated standard error and stack the resulting moments in $g(\theta)$. The estimator is

$$\hat{\theta} = \arg \min_{\theta \in \Theta} g(\theta)' W g(\theta). \quad (21)$$

We impose positivity of costs and scale parameters by optimising their logarithms. The weighting matrix uses the inverse of the jointly estimated covariance matrix of the moments. We estimate this covariance matrix with a common cluster bootstrap that recomputes all pre-reform and difference-in-differences moments in each draw.

The overidentifying-restrictions statistic is

$$J = g(\hat{\theta})' W g(\hat{\theta}),$$

which is asymptotically chi-squared with the number of moments minus eight degrees of freedom. The cluster bootstrap also provides sampling uncertainty for the parameter estimates and the decomposition.

6.4 Decomposition

For outcome Y , let $Y_j(w_A; \hat{\theta})$ be the model outcome in occupation j when the wage in A is w_A . Holding w_B , the worker distribution, notification rates, and all estimated parameters fixed, define

$$\Delta Y_j = Y_j(w_A^1; \hat{\theta}) - Y_j(w_A^0; \hat{\theta}).$$

The structural counterpart of the empirical difference-in-differences coefficient is

$$\tau_Y^M = \underbrace{\Delta Y_A}_{\text{effect on treated occupation}} + \underbrace{(-\Delta Y_B)}_{\text{control-group spillover}}. \quad (22)$$

The first term is the full equilibrium response in A . It includes the direct wage incentive as well as endogenous changes in competition and continuation values. The second term measures contamination of the comparison group. If workers substitute away from B , then $\Delta Y_B < 0$ for applications and the DiD exceeds the equilibrium effect on A .

We further decompose the spillover on B . Along the equilibrium wage path,

$$\frac{d\lambda_B(q)}{dw_A} = \underbrace{-\frac{\lambda_B(q)[1 - \lambda_B(q)]}{s_B} \frac{du(q)}{dw_A}}_{\text{continuation-value channel}} + \underbrace{\frac{\lambda_B(q)[1 - \lambda_B(q)]}{s_B} \frac{c_B}{p_B(q_B)^2} \frac{dp_B(q_B)}{dw_A}}_{\text{hiring-probability channel}}. \quad (23)$$

The first channel normally reduces applications to B : a better opportunity in A raises the value of continuing to search. The second channel often mitigates that reduction: fewer applications to B lower competition and raise the probability of being hired there.

We aggregate each derivative over worker types and integrate it from w_A^0 to w_A^1 . For applications and average quality this yields

$$\Delta Y_B = \Delta Y_B^u + \Delta Y_B^p, \quad \tau_Y^M = \Delta Y_A - \Delta Y_B^u - \Delta Y_B^p.$$

Table 6: Structural parameter estimates

Parameter	Interpretation	Estimate
b	Baseline value of continued search	0.1637
b_1	Slope of continued-search value in average quality rank	0.6737
c_A	Application cost, occupation A	0.1023
c_B	Application cost, occupation B	0.0875
s_A	Worker payoff dispersion, occupation A	0.4976
s_B	Worker payoff dispersion, occupation B	0.3106
σ_A	Employer-choice dispersion, occupation A	0.7879
σ_B	Employer-choice dispersion, occupation B	0.7366

Notes: The table reports GMM estimates of the behavioural parameters.

Table 7: Model fit and wage-change moments

Moment	Data	Model	Standardised residual
Pre-reform applications per vacancy, A	10.270387	9.626326	1.496912
Pre-reform applications per vacancy, B	7.735478	7.481372	0.817640
DiD: applications	0.210400	0.220075	-0.437778
DiD: average applicant quality	0.009500	0.003960	1.629384

Notes: Pre-reform applications are expressed in levels; the application DiD is expressed in log points; and applicant quality is measured in occupation-specific percentile-rank units. The table reports the data, their model counterparts, and standardised residuals.

The decomposition is performed in the same units as the empirical outcome. Appendix D gives the aggregation formulas and numerical implementation.

6.5 Estimates and decomposition


This subsection reports the structural parameter estimates, model fit, and decomposition.

Table 6 reports the eight estimated behavioural parameters. The remaining inputs are calibrated directly from the pre-reform aggregate moments and the wage-experiment context. Table 7 reports the aggregate moments most directly related to the wage experiment and compares the targeted DiD moments with their model counterparts. The model closely reproduces the pre-reform application rates and the wage-induced changes in applications and applicant quality.

Figure 10 presents the main decomposition as an accounting identity. The equilibrium effect on A and the contribution from spillovers onto the comparison group, $-\Delta Y_B$, sum to the structural DiD. The lower box reports the empirical DiD supplied to the estimator.

Table 8 reports the same decomposition and splits the B response into its two structural

Figure 10: Structural decomposition of the DiD



Notes: The control-group component equals minus the model response in occupation B , so a decline in B contributes positively to the DiD. The arrow links the structural DiD to its empirical counterpart.

channels: the continuation-value effect and the hiring-probability effect.

For applications, the model-based difference-in-differences effect of 0.214 log points comprises a 0.200-log-point increase in occupation A and a 0.013-log-point decline in occupation B . Comparison-group displacement therefore accounts for about 6% of the contrast, while the equilibrium response in the treated occupation accounts for the remaining 94%. Nearly all of the response in B operates through the hiring-probability channel. For average applicant quality, the response in B is close to zero, so the model-based difference-in-differences effect is almost entirely attributable to the improvement in occupation A .

Table 8: Decomposition of the reduced-form estimators

<i>Decomposition of the DiD effect</i>				
Outcome	Empirical DiD	Model DiD	Effect on A	Effect on B
Applications	0.210400	0.213645	0.200175	-0.013470
Mean quality	0.009500	0.003992	0.004050	0.000058
<i>Decomposition of the effect on B</i>				
Outcome	Continuation value	Hiring prob.	Adding-up err.	
Applications	-0.000324	-0.013146	4.90×10^{-10}	
Mean quality	-0.000012	0.000070	7.61×10^{-13}	

Notes: The model DiD equals the effect on A minus the effect on B . Within B , the continuation-value and hiring-probability channels sum to the total B effect up to numerical path-integration error.

7 Conclusion

We study an occupation-wide increase in the default wage on a platform for short-term hospitality work. Raising the waiting-staff wage from €16 to €17 increased applications by approximately 23% relative to three neighbouring occupations, implying a wage elasticity of about 3.5. Mean applicant quality rose modestly, with an elasticity of approximately 0.3, and the improvement was concentrated at the bottom of the applicant-quality distribution. Hires increased by about 11%, implying an elasticity of 1.8. Together, these estimates show that occupation-wide wage changes can affect not only how many workers apply, but also the composition of the applicant pool and the probability that employers fill their vacancies.

The platform setting also makes clear why a conventional difference-in-differences estimate need not equal the equilibrium response in the treated occupation. Workers receive opportunities in several occupations, so a higher wage in one occupation can divert applications from another. Our structural decomposition shows that this spillover is modest. The decline in applications to the comparison occupations accounts for about 6% of the model-based application contrast, while the remaining 94% reflects the equilibrium response in the treated occupation. The spillover in applicant quality is economically negligible. The reduced-form elasticities therefore closely approximate the equilibrium responses of waiting-staff vacancies, and the conclusion that the higher wage attracted substantially more and somewhat better applicants is robust to interference across occupations.

References

- ALTMANN, S., A. GLENNY, R. MAHLSTEDT, AND A. SEBALD (2023): “The Direct and Indirect Effects of Online Job Search Advice,” .
- AZAR, J., E. HUET-VAUGHN, I. MARINESCU, B. TASKA, AND T. VON WACHTER (2024): “Minimum Wage Employment Effects and Labor Market Concentration,” *Review of Economic Studies*, 91, 1843–1883.
- BALGOVA, M., T. TEKLESELASSIE, L. HENSEL, AND M. J. WITTE (2025): “Wage Information and Applicant Selection,” IZA Discussion Paper 18220, Institute of Labor Economics (IZA).
- BANFI, S. AND B. VILLENA-ROLDÁN (2019): “Do High-Wage Jobs Attract More Applicants? Directed Search Evidence from the Online Labor Market,” *Journal of Labor Economics*, 37, 715–746.
- BASSIER, I. (2022): “Collective Bargaining and Spillovers in Local Labor Markets,” CEP Discussion Paper 1895, Centre for Economic Performance, London School of Economics and Political Science.
- BASSIER, I., A. DUBE, AND S. NAIDU (2022): “Monopsony in Movers: The Elasticity of Labor Supply to Firm Wage Policies,” *Journal of Human Resources*, 57, S50–S86.
- BASSIER, I., A. MANNING, AND B. PETRONGOLO (2025): “Vacancy Duration and Wages,” *The Review of Economics and Statistics*, 1–28.
- BEERLI, A., S. FIORIN, A. GULYAS, D. KOPP, M. KHOSHNAMEA, AND M. SIEGENTHALER (2026): “The Role of Wages and Job Benefits in Job Search: Evidence from a Large-Scale Online Field Experiment,” .
- BELOT, M., B. K. DE KONING, D. FOUARGE, P. KIRCHER, P. MULLER, AND S. PHILIPPEN (2025): “Advising Job Seekers in Occupations with Poor Prospects: A Field Experiment,” .
- BELOT, M., P. KIRCHER, AND P. MULLER (2019): “Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice,” *The Review of Economic Studies*, 86, 1411–1447.
- (2022): “How Wage Announcements Affect Job Search—A Field Experiment,” *American Economic Journal: Macroeconomics*, 14, 1–67.
- (2026): “Do the Long-term Unemployed Benefit from Automated Occupational Advice during Online Job Search?” *The Economic Journal*, 136, 184–206.
- CALDWELL, S. AND O. DANIELI (2024): “Outside Options in the Labour Market,” *Review of Economic Studies*, 91, 3286–3315.
- CARD, D. (2022): “Who Set Your Wage?” *American Economic Review*, 112, 1075–1090.
- CENGIZ, D., A. DUBE, A. LINDNER, AND B. ZIPPERER (2019): “The Effect of Minimum Wages on Low-Wage Jobs,” *The Quarterly Journal of Economics*, 134, 1405–1454.

- CLEMENS, J., L. B. KAHN, AND J. MEER (2021): “Dropouts Need Not Apply? The Minimum Wage and Skill Upgrading,” *Journal of Labor Economics*, 39, S107–S149.
- DAL BÓ, E., F. FINAN, AND M. A. ROSSI (2013): “Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service,” *The Quarterly Journal of Economics*, 128, 1169–1218.
- DATTA, N. (2023): “The Measure of Monopsony: The Labour Supply Elasticity to the Firm and Its Constituents,” .
- DEMIR, G. K. (2024): “Labor Market Frictions and Spillover Effects from Publicly Announced Sectoral Minimum Wages,” .
- DERENONCOURT, E. AND C. MONTIALOUX (2021): “Minimum Wages and Racial Inequality,” *The Quarterly Journal of Economics*, 136, 169–228.
- DERENONCOURT, E. AND D. WEIL (2025): “Voluntary Minimum Wages: The Local Labor Market Effects of National Retailer Policies,” *Quarterly Journal of Economics*, 140, 1901–1958.
- DUBE, A., J. JACOBS, S. NAIDU, AND S. SURI (2020): “Monopsony in Online Labor Markets,” *American Economic Review: Insights*, 2, 33–46.
- DUSTMANN, C., A. LINDNER, U. SCHÖNBERG, M. UMKEHRER, AND P. VOM BERGE (2022): “Reallocation Effects of the Minimum Wage,” *The Quarterly Journal of Economics*, 137, 267–328.
- GIUPPONI, G., R. JOYCE, A. LINDNER, T. WATERS, T. WERNHAM, AND X. XU (2024): “The Employment and Distributional Impacts of Nationwide Minimum Wage Changes,” *Journal of Labor Economics*, 42, S293–S333.
- GIUPPONI, G. AND S. J. MACHIN (2026): “Company Wage Policy in a Low-Wage Labor Market,” .
- GOPALAN, R., B. H. HAMILTON, A. KALDA, AND D. SOVICH (2021): “State Minimum Wages, Employment, and Wage Spillovers: Evidence from Administrative Payroll Data,” *Journal of Labor Economics*, 39, 673–707.
- HARASZTOSI, P. AND A. LINDNER (2019): “Who Pays for the Minimum Wage?” *The American Economic Review*, 109, 2693–2727.
- HIRSCH, B., E. J. JAHN, A. MANNING, AND M. OBERFICHTNER (2022): “The Wage Elasticity of Recruitment,” .
- JÄGER, S., C. ROTH, N. ROUSSILLE, AND B. SCHOEFER (2024): “Worker Beliefs About Outside Options,” *The Quarterly Journal of Economics*, 139, 1505–1556, eprint: <https://academic.oup.com/qje/article-pdf/139/3/1505/58456518/qjae001.pdf>.
- KLAEUI, J. (2026): “When to Broaden, When to Focus: Job Search and Market Tightness,” .
- MANNING, A. (2021): “Monopsony in Labor Markets: A Review,” *ILR Review*, 74, 3–26.

- MARINESCU, I. AND R. WOLTHOFF (2020): "Opening the Black Box of the Matching Function: The Power of Words," *Journal of Labor Economics*, 38, 535–568.
- MATSUDAIRA, J. D. (2014): "Monopsony in the Low-Wage Labor Market? Evidence from Minimum Nurse Staffing Regulations," *Review of Economics and Statistics*, 96, 92–102.
- RANSOM, M. R. AND D. P. SIMS (2010): "Estimating the Firm's Labor Supply Curve in a "New Monopsony" Framework: Schoolteachers in Missouri," *Journal of Labor Economics*, 28, 331–355.
- SOKOLOVA, A. AND T. SORENSEN (2021): "Monopsony in Labor Markets: A Meta-Analysis," *ILR Review*, 74, 27–55.
- STAIGER, D. O., J. SPETZ, AND C. S. PHIBBS (2010): "Is There Monopsony in the Labor Market? Evidence from a Natural Experiment," *Journal of Labor Economics*, 28, 211–236.

A Screenshots of the platform

A.1 Firm user experience




(a) Category of establishment choice



(b) Establishment description



(c) Qualification family choice




(d) Qualification choice




(e) Type of contract choice




(f) Number of required workers




(g) Dates of job




(h) Detailed schedule of job




(i) Job description




(j) Wage choice



(k) Job ad summary




(l) Number of applicants to the ad



(m) Applicant profile




(n) Applicant profile



(o) Applicant selection


A.2 Worker user experience




(a) Job ad example




(b) Job ad summary



(c) Establishment description



(d) Application final step



(e) Application tracking

B Estimating quality

We use employers’ revealed choices to construct a quality score for each worker–job pair. Let H_{ij} indicate that worker i is hired for job j , conditional on applying. We estimate a class-weighted logistic model and define the fitted score as

$$S_{ij} = \frac{1}{1 + \exp(-z_{ij})}, \quad z_{ij} = \beta_0 + X'_{ij}\beta + \gamma_{c(j),e(j)}. \quad (24)$$

The vector X_{ij} contains worker-, job-, and pair-level characteristics, and $\gamma_{c(j),e(j)}$ denotes establishment fixed effects indexed by establishment category and employer. Worker characteristics include gender, age quintile, average rating, and time since the last connection to the platform. The job controls include the number of workers requested. Pair-level characteristics include quintiles of straight-line distance between the worker’s residence and the workplace; indicators for living in the same postcode or *département*; whether the worker has previously worked for and been approved by the employer; and detailed measures of previous experience. We group experience into the most recent assignment, the second through fifth most recent assignments, and the sixth through tenth most recent assignments. For each group, we include average ratings and indicators for whether the experience was obtained through the platform, was in the same occupation or establishment category as the current job, or was in a higher, lower, or undefined occupation within the same family. We rank occupations within families by their mean posted hourly wage. Missing-value indicators are included for age, average ratings, experience ratings, and time since the last connection.

Hires are rare, so the estimation weights hired and non-hired observations inversely by their respective sample frequencies. This weighting improves the model’s ability to rank hires but changes the calibration of its fitted values. We therefore interpret S_{ij} as a quality score, not as the applicant’s actual hiring probability.

We randomly assign 20% of service-sector jobs posted between 1 January 2023 and 30 December 2025 to the training sample and use the remaining 80% as the test sample. The broader service-sector data contain 22 occupations. Eighteen are not suitable comparison occupations for the event study, but they provide useful information about employers’ hiring choices. Before estimation-specific exclusions, the training sample contains 382,201 worker–job pairs, 10,989 workers, and 23,165 jobs; the test sample contains 4,155,162 pairs, 20,708 workers, and 237,945 jobs. The effective logistic-regression sample contains 361,689 observations after fixed-effect and missing-value handling. Quality is pair-specific: a worker can receive different scores for different jobs.

Table A1 compares worker–job pairs in the training, test, and DiD samples. Although the test group contains 80% of jobs, it contains 92% of pairs because applicant-pool sizes vary substantially across postings. Random assignment makes the training and test samples broadly comparable in worker demographics, geographic proximity, prior worker–employer relationships, ratings, job requirements, and work experience. The DiD sample is instead determined by the event-study period and four selected occupations. It consequently contains more women, younger workers, and somewhat different occupational histories. We obtain the quality outcome used in the DiD analysis by applying the model estimated on the training sample to this analysis sample.

Table A1: Summary statistics of worker-job pairs across samples

	Training sample	Test sample	DiD sample
Female	0.44 (0.50)	0.37 (0.48)	0.49 (0.50)
Age	29.37 (7.15)	30.21 (7.69)	28.15 (6.47)
Same postcode	0.02 (0.13)	0.02 (0.12)	0.02 (0.13)
Same department	0.32 (0.47)	0.29 (0.45)	0.31 (0.46)
Prior worker-establishment history	0.08 (0.27)	0.06 (0.24)	0.07 (0.25)
Average rating	4.59 (0.52)	4.58 (0.53)	4.53 (0.58)
Number of workers requested	1.32 (1.20)	1.21 (1.13)	1.33 (1.38)
Straight-line distance (km)	15.68 (41.37)	16.55 (43.12)	17.27 (48.26)
Experience (most recent)			
Average rating	4.79 (0.83)	4.81 (0.80)	4.80 (0.82)
Same qualification	0.33 (0.47)	0.37 (0.48)	0.35 (0.48)
Same category	0.12 (0.33)	0.15 (0.35)	0.10 (0.30)
Rank	13.79 (4.79)	14.59 (5.66)	14.71 (4.81)
Same family, higher rank	0.27 (0.44)	0.21 (0.41)	0.13 (0.33)
Same family, lower rank	0.20 (0.40)	0.16 (0.37)	0.25 (0.44)
Different family	0.20 (0.40)	0.26 (0.44)	0.27 (0.44)
Experience (last 2-5)			
Average rating	4.65 (0.94)	4.66 (0.94)	4.62 (0.97)
Same qualification	0.34 (0.48)	0.37 (0.48)	0.37 (0.48)
Same category	0.12 (0.32)	0.14 (0.35)	0.10 (0.30)
Rank	13.84 (4.69)	14.64 (5.64)	14.73 (4.78)
Same family, higher rank	0.28 (0.45)	0.21 (0.41)	0.12 (0.32)
Same family, lower rank	0.18 (0.39)	0.16 (0.37)	0.24 (0.43)
Different family	0.18 (0.39)	0.25 (0.43)	0.25 (0.43)
Experience (last 6-10)			
Average rating	4.69 (0.84)	4.70 (0.85)	4.66 (0.90)
Same qualification	0.30 (0.46)	0.33 (0.47)	0.31 (0.46)
Same category	0.10 (0.30)	0.12 (0.33)	0.08 (0.28)
Rank	12.68 (5.91)	13.10 (6.94)	13.14 (6.46)
Same family, higher rank	0.26 (0.44)	0.20 (0.40)	0.11 (0.31)
Same family, lower rank	0.17 (0.37)	0.14 (0.35)	0.23 (0.42)
Different family	0.18 (0.38)	0.22 (0.41)	0.23 (0.42)
Worker-job pairs	382,201	4,155,162	320,833
Share of sample	0.08	0.92	0.07

Notes: The table reports summary statistics for worker-job pairs in the training, test, and DiD samples. Standard deviations are in parentheses. The DiD sample contains applications to waiting-staff, bar-staff, receptionist, and runner jobs posted between 5 April 2023 and 4 May 2024.

Table A2 reports selected variables with the largest absolute coefficients in the service-sector logistic model. These raw logit coefficients are descriptive of the fitted model; because the regressors use different scales, their magnitudes should not be interpreted as a formal variable-importance ranking.

Table A2: Selected large coefficients from the quality model

	Service sector (1)
Previously worked for and approved by the establishment	2.41*** (0.009)
Missing average rating	0.619*** (0.129)
Last experience from the same qualification	-0.941*** (0.067)
Last experience from a higher qualification in the same family	-1.17*** (0.068)
Last experience from a lower qualification in the same family	-0.923*** (0.067)
Last experience from a different family	-1.27*** (0.067)
Median qualification of the last 2-5 experiences is from the same qualification	0.605*** (0.050)
Median qualification of last 2-5 experiences is from a lower qualification in the same family	0.474*** (0.050)
Dependent variable mean	0.042
Observations	361,689
Pseudo R ²	0.249
Establishment fixed effects	✓

Notes: The table reports selected coefficients with the largest absolute values in equation (24). Standard errors are in parentheses.

Table A3 evaluates the quality model in the test sample. An AUC of 0.76 means that the model assigns a higher score to a randomly selected hired applicant than to a randomly selected non-hired applicant 76% of the time. At the threshold of 0.6, which maximises Youden's J , sensitivity is 0.605 and specificity is 0.785. Thus Youden's $J = 0.605 + 0.785 - 1 = 0.390$. Precision is 0.15, reflecting the rarity of hiring: only approximately 3.7% of test-sample applications result in a hire. These metrics show that the score contains meaningful ranking information, while also illustrating the difficulty of classifying rare outcomes.


Table A3: Prediction evaluation metrics

AUC	Sensitivity	Specificity	Accuracy	Precision	Youden Index	Threshold
0.76	0.605	0.785	0.774	0.15	0.390	0.6

Notes: The table reports classification metrics for the applicant-quality model in the test sample. The threshold maximises Youden’s J , defined as sensitivity plus specificity minus one.

Figure A3 plots the receiver operating characteristic (ROC) curve. The curve lies above the 45-degree line associated with random classification. Its shape illustrates the trade-off between sensitivity and specificity: lowering the classification threshold increases both the true-positive and false-positive rates. The AUC summarises the model’s ability to rank applicants across all possible thresholds.

Figure A3: Receiver operating characteristic curve



Notes: The curve plots sensitivity against one minus specificity across all classification thresholds in the test sample. Table A3 reports the corresponding evaluation metrics.

C Conceptual framework: derivations and equilibrium

This appendix derives the worker’s application rule, defines the exact Poisson-logit equilibrium, relates it to the mean-field representation used in the main text and estimation, and proves Proposition 1. For compactness, some derivations use continuous notation. In the structural estimation, these integrals are evaluated as sums over the empirical joint cells, and $\alpha_j(q)$ denotes the observed cell-specific notification rate.

C.1 Worker problem

Let $W_j(\varepsilon_j; q, w)$ be the value of an employment spell in occupation j with match-specific payoff ε_j . Employment ends at the close of a period with probability δ , so

$$W_j(\varepsilon_j; q, w) = w_j + \varepsilon_j + \tilde{\beta} [(1 - \delta)W_j(\varepsilon_j; q, w) + \delta U(q; w)] . \quad (25)$$

It follows that

$$W_j(\varepsilon_j; q, w) - U(q; w) = \frac{w_j + \varepsilon_j - (1 - \tilde{\beta})U(q; w)}{1 - \tilde{\beta}(1 - \delta)} . \quad (26)$$

Before receiving a notification, the worker's value satisfies

$$U(q; w) = b_0 + b_1 \bar{q} + \tilde{\beta} U(q; w) + \sum_{j \in \{A, B\}} \alpha_j(q) \mathbb{E}_{\varepsilon_j} \left[\max \left\{ 0, -\tilde{c}_j + \tilde{\beta} p_j(q_j; w) [W_j(\varepsilon_j; q, w) - U(q; w)] \right\} \right] , \quad (27)$$

where b_0 is the baseline flow payoff while non-employed, $\bar{q} = (q_A + q_B)/2$, and b_1 allows the value of waiting to vary with worker quality. Define

$$u = (1 - \tilde{\beta})U, \quad \beta = \frac{\tilde{\beta}}{1 - \tilde{\beta}(1 - \delta)}, \quad c_j = \frac{\tilde{c}_j}{\tilde{\beta}},$$

and absorb the scale conversion of $b_0 + b_1 \bar{q}$ into $b(q) = b + b_1 \bar{q}$. Substitution of (26) into (27) gives the application rule

$$p_j(q_j; w) [w_j + \varepsilon_j - u(q; w)] \geq c_j.$$

C.2 Application threshold

Rearranging the application rule yields the threshold

$$\varepsilon_j \geq \bar{\varepsilon}_j(q; w) \equiv \frac{c_j}{p_j(q_j; w)} - [w_j - u(q; w)].$$

For a logistic shock with scale s_j ,

$$\Pr(\varepsilon_j \geq x) = \frac{1}{1 + \exp(x/s_j)}.$$

Evaluating this survival function at $\bar{\varepsilon}_j(q; w)$ proves equation (2). Its index is increasing in w_j and p_j , holding u fixed, which establishes the stated monotonicity.

The value of the option to apply can also be written in terms of the application probability. For a logistic random variable,

$$\mathbb{E}[(\varepsilon_j - \bar{\varepsilon}_j)^+] = -s_j \log(1 - \lambda_j).$$

The transformed search value is therefore

$$u(q; w) = b - \sum_{j \in \{A, B\}} \alpha_j(q) \beta s_j p_j(q; w) \log[1 - \lambda_j(q; w)],$$

which is equation (3).

C.3 Worker policy functions

Fix wages w and hiring probabilities p_A, p_B . For each worker type, the continuation value is the solution to

$$\Psi(u; p, w) \equiv u - b(q) + \sum_{j \in \{A, B\}} \alpha_j(q) \beta s_j p_j \log \left[1 - \frac{1}{1 + \exp\{-[w_j - u - c_j/p_j]/s_j\}} \right] = 0. \quad (28)$$

Because

$$\frac{\partial \Psi}{\partial u} = 1 + \sum_{j \in \{A, B\}} \alpha_j(q) \beta p_j \lambda_j > 0,$$

the worker problem has at most one solution. Its limits as $u \rightarrow -\infty$ and $u \rightarrow +\infty$ have opposite signs, so a solution exists. Denote it by $m(p_A, p_B; w)$. The worker policy functions are

$$\ell_j(p_A, p_B; w) = \frac{1}{1 + \exp\{-[w_j - m(p_A, p_B; w) - c_j/p_j]/s_j\}}.$$

C.4 Exact Poisson-logit hiring operator

Let

$$a_j(x) = \exp(x/\sigma_j)$$

be the employer-choice index. Conditional on a focal worker applying with productivity q_j , the logit probability of being selected is

$$\frac{a_j(q_j)}{a_j(q_j) + \sum_{i \in \mathcal{A}_j} a_j(q_{ij})}.$$

Applications from other notified workers form a marked Poisson process. Let ν_j denote the distribution of productivity ranks among workers notified about an occupation- j vacancy, and define

$$\bar{\lambda}_j(x; w) = \mathbb{E}[\lambda_j(Q_A, Q_B; w) \mid Q_j = x, \text{an occupation-}j \text{ notification}].$$

Using

$$\frac{a}{a + S} = \int_0^\infty a \exp[-t(a + S)] dt,$$

the probability of being hired conditional on applying is

$$\begin{aligned} p_j(q_j; w) &= \kappa_j \int_0^\infty a_j(q_j) \exp[-ta_j(q_j)] \\ &\quad \times \exp \left\{ -M_j \int_{[0,1]} [1 - \exp\{-ta_j(x)\}] \bar{\lambda}_j(x; w) d\nu_j(x) \right\} dt. \end{aligned} \quad (29)$$

An exact equilibrium is a collection $(p_A, p_B, \lambda_A, \lambda_B, u)$ satisfying

$$\lambda_j(q; w) = \ell_j(p_A(q_A; w), p_B(q_B; w); w), \quad (30)$$

$$p_j(q_j; w) = \mathcal{P}_j[\bar{\lambda}_j(\cdot; w)](q_j), \quad (31)$$

where $\mathcal{P}_j$ is the operator in (29). The quantitative implementation assumes that this fixed point is unique. More generally, the analysis requires a stable equilibrium selection rule.

C.5 Mean-field approximation

The exact operator requires a one-dimensional integral for each occupation and productivity point. The mean-field approximation replaces the random logit denominator by its expectation:

$$D_j(w) = M_j \int_{[0,1]} a_j(x) \bar{\lambda}_j(x; w) d\nu_j(x), \quad p_j^{\text{MF}}(q_j; w) = \kappa_j \frac{a_j(q_j)}{a_j(q_j) + D_j(w)}.$$

This is the representation in equations (4) and (5). It preserves smooth employer choice, productivity sorting, and the endogenous response of hiring probabilities to the applicant pool while substantially reducing computation.

The limiting cases are useful for interpretation. As $\sigma_j \downarrow 0$, employers select the highest-productivity applicant with probability approaching one. As $\sigma_j \rightarrow \infty$, employers select approximately at random among applicants and

$$p_j(q_j; w) \rightarrow \kappa_j \frac{1 - \exp[-\mathcal{A}_j(w)]}{\mathcal{A}_j(w)},$$

independently of q_j , where $\mathcal{A}_j = M_j \int \bar{\lambda}_j(x; w) d\nu_j(x)$.

C.6 Proof of Proposition 1

Write equation (2) as

$$\lambda_j = L \left(\frac{w_j - u - c_j/p_j}{s_j} \right), \quad L'(z) = L(z)[1 - L(z)].$$

Taking the total derivative along an equilibrium wage path gives

$$\frac{d\lambda_j}{dw_k} = \frac{\lambda_j(1 - \lambda_j)}{s_j} \left[\mathbf{1}\{j = k\} - \frac{du}{dw_k} + \frac{c_j}{p_j^2} \frac{dp_j}{dw_k} \right],$$

which proves equation (6). The first term is positive for an own-wage change, while the other two terms need not be. Therefore the total derivative is not sign restricted. $\square$

For completeness, define

$$H_j(q) = -s_j \log[1 - \lambda_j(q)], \quad K_j(q) = H_j(q) + \frac{c_j}{p_j(q)} \lambda_j(q).$$

Differentiating the continuation-value equation and using the envelope property gives

$$\frac{du(q)}{dw_k} = \frac{\beta \alpha_k(q) p_k(q) \lambda_k(q) + \beta \sum_{j \in \{A, B\}} \alpha_j(q) K_j(q) \frac{dp_j(q)}{dw_k}}{1 + \beta \sum_{j \in \{A, B\}} \alpha_j(q) p_j(q) \lambda_j(q)}. \quad (32)$$

Holding hiring probabilities fixed, an increase in w_k strictly raises the continuation value for workers with a positive probability of applying to and obtaining a job in k . Hiring-probability adjustments can reinforce or attenuate this response.

Under the mean-field rule,

$$\frac{dp_j(q_j)}{dw_k} = -\frac{p_j(q_j)}{a_j(q_j) + D_j} \frac{dD_j}{dw_k}, \quad (33)$$

$$\frac{dD_j}{dw_k} = M_j \int_{[0,1]} a_j(x) \frac{d\bar{\lambda}_j(x)}{dw_k} dv_j(x). \quad (34)$$

Equations (6), (32), and (33)–(34) form a linear fixed-point system for local equilibrium responses.

C.7 Applicant-pool outcomes

Applications to an occupation- j vacancy form a marked Poisson process with rank-intensity measure

$$\mu_j(dx; w) = M_j \bar{\lambda}_j(x; w) dv_j(x), \quad x \in [0, 1].$$

The expected number of applications is $\mathcal{A}_j = \int_{[0,1]} \mu_j(dx)$, and the probability of a nonempty applicant pool is $1 - \exp(-\mathcal{A}_j)$. Conditional on a nonempty pool, mean quality is

$$\bar{Q}_j = \frac{\int_{[0,1]} x \mu_j(dx)}{\mathcal{A}_j}.$$

Let

$$C_j(x) = \int_{[0,x]} \mu_j(dv), \quad T_j(x) = \int_{[x,1]} \mu_j(dv).$$

Conditional on a nonempty pool, expected minimum and maximum quality are

$$Q_{j,\min} = \int_0^1 \frac{\exp[-C_j(x)] - \exp[-\mathcal{A}_j]}{1 - \exp[-\mathcal{A}_j]} dx, \quad (35)$$

$$Q_{j,\max} = \int_0^1 \frac{1 - \exp[-T_j(x)]}{1 - \exp[-\mathcal{A}_j]} dx. \quad (36)$$

These expressions must be adapted if the empirical analysis assigns values to vacancies with no applications rather than conditioning on nonempty applicant pools.

D Structural estimation and decomposition: additional details

This appendix describes the construction of the empirical moments, provides a formal identification argument for the mean-field model, and gives the GMM moment conditions and inference procedure. The current baseline estimation uses empirical joint cells and cell-specific notification rates. Gaussian-copula notation appears only as an optional sensitivity case.

D.1 Empirical moment construction

D.1.1 Reference population and productivity cells

All baseline moments are constructed from the stable pre-reform period and refer to the same population of at-risk workers and vacancies. Worker quality measures must be determined before the application decision and converted into occupation-specific percentile ranks $q_A, q_B \in [0, 1]$. Ranks are defined in the searching or at-risk population rather than in the selected applicant pool. When quality is estimated, sample splitting or cross-fitting prevents a worker's realised application or hiring outcome for the same job from mechanically determining that worker's rank. Because the reduced-form score S_{ij} is worker-job specific, we construct the worker-level ranks (q_A, q_B) by aggregating predicted quality across potential jobs in each occupation before the application decision.

We partition each rank and cross the partitions. A joint cell c contains workers with similar (q_A, q_B) . Its weight is the share of at-risk workers or worker-periods in that cell before notification and application decisions. Applicant or hire shares are not valid cell weights because they are endogenous outcomes.

D.1.2 Conditional application and hiring moments

For each joint cell, the application moments are

$$\hat{\lambda}_A(c) = \frac{\#\{\text{A applications following A notifications in } c\}}{\#\{\text{A notifications in } c\}}, \quad (37)$$

$$\hat{\lambda}_B(c) = \frac{\#\{\text{B applications following B notifications in } c\}}{\#\{\text{B notifications in } c\}}. \quad (38)$$

The denominator is the relevant number of notifications, not the number of workers. Notifications with no application remain in the denominator.

The model implies $p_A = p_A(q_A)$ and $p_B = p_B(q_B)$. We therefore estimate

$$\hat{p}_A(x) = \frac{\#\{\text{hires from A applications in } q_A \text{ bin } x\}}{\#\{\text{A applications in } q_A \text{ bin } x\}}, \quad (39)$$

$$\hat{p}_B(x) = \frac{\#\{\text{hires from B applications in } q_B \text{ bin } x\}}{\#\{\text{B applications in } q_B \text{ bin } x\}}. \quad (40)$$

Thus hiring probabilities are conditional on applying. Hires divided by notifications combine p_j and λ_j and are not the relevant moment. We merge the one-dimensional hiring estimates back onto the joint cells for use in the application equations, but count each distinct rank-bin hiring estimate only once in the GMM criterion.

D.1.3 Aggregate moments

For occupation j , the aggregate pre-reform moments are

$$\hat{M}_j = \frac{\text{notifications sent by } j}{\text{vacancies in } j}, \quad \hat{\mathcal{A}}_j = \frac{\text{applications received by } j}{\text{vacancies in } j}, \quad (41)$$

$$\hat{z}_j = \frac{\text{vacancies ending without a hire}}{\text{vacancies in } j}, \quad \hat{\alpha}_j = \frac{\text{at-risk worker-periods with a } j \text{ notification}}{\text{at-risk worker-periods}}. \quad (42)$$

For each joint cell c , we also compute

$$\hat{\alpha}_j(c) = \frac{j\text{-notifications sent to worker-periods in cell } c}{\text{at-risk worker-periods in cell } c}.$$

The aggregate $\hat{\alpha}_j$ is a weighted average of these cell rates, but the structural equilibrium uses $\hat{\alpha}_j(c)$. Applications per vacancy are derived from M_j , the cell application rates, and the occupation-specific notification shares. The time unit used to calculate notification rates is also used for match duration and discounting. Standard errors account for repeated workers and common vacancy or time shocks using the clustering structure of the empirical analysis.

D.1.4 Representative comparison occupation

The empirical comparison group contains several non-A occupations. In the two-occupation approximation, we pool their notifications, applications, vacancies, and hires before con-

structing the B moments. We define the representative pre-reform wage as

$$w_B = \sum_{k \in \mathcal{B}} \omega_k^N w_k, \quad \omega_k^N = \frac{\text{notifications from occupation } k}{\text{all B notifications}}. \quad (43)$$

This common w_B is used in every joint cell. The approximation is most credible when B occupations have similar wages and technologies and their notification mix is stable across worker types and time. Robustness exercises can restrict the comparison set or model its constituent occupations separately.

D.1.5 Difference-in-differences moments

The empirical DiD moments are the interaction coefficients from the paper's preferred specifications for applications and average applicant quality. The application coefficient comes from a Poisson model and is therefore a change in a log application rate. Its structural counterpart must use $h(x) = \log x$. The level and $\log(1+x)$ transformations are relevant only for alternative empirical specifications estimated in those units.

For quality, each A application is assigned the applicant's q_A and each B application the applicant's q_B . The occupation-period outcome is the application-weighted mean rank. The empirical quality outcome is constructed in these occupation-specific rank units, so the empirical and structural DiDs are directly comparable.

D.2 Formal identification

Let the directly observed or calibrated parameter vector be

$$\theta_0 = (\delta, \tilde{\beta}, \beta, M_A, M_B, \kappa_A, \kappa_B, \{\alpha_A(c), \alpha_B(c), \omega_c, q_A(c), q_B(c)\}_c),$$

and let

$$\theta = (b, b_1, c_A, c_B, s_A, s_B, \sigma_A, \sigma_B)$$

be the behavioural parameter vector.

We impose the following conditions. First, the time unit, mean match duration, wages, notification frequencies, notified workers per vacancy, and the empirical joint distribution of productivity ranks are observed. Second, the conditional moments are interior:

$$0 < \lambda_j(q) < 1, \quad 0 < p_j(q_j) < \kappa_j.$$

Each occupation contains at least two productivity values with distinct hiring probabilities.

Third, the mean-field equilibrium is unique or a stable equilibrium-selection rule is fixed. Finally, the rank condition defined below holds.

D.2.1 Parameters identified directly

Suppose there are T model periods per day and mean match duration is L days. Geometric separation gives

$$\delta = \frac{1}{LT}. \quad (44)$$

Given annual discount rate r ,

$$\tilde{\beta} = \exp\left(-\frac{r}{365T}\right), \quad \beta = \frac{\tilde{\beta}}{1 - \tilde{\beta}(1 - \delta)}. \quad (45)$$

Equation (15) identifies κ_j , while $\alpha_j(c)$, the cell weights, and M_j are direct notification and risk-set moments. If a Gaussian copula is imposed as a sensitivity check and τ_q is Kendall's rank correlation,

$$\rho = \sin\left(\frac{\pi}{2}\tau_q\right). \quad (46)$$

D.2.2 Employer-choice dispersion

Proposition 2 (Employer selectivity). *Conditional hiring probabilities $p_j(q_j)$, together with κ_j , point identify σ_j . For two distinct ranks q_j^1, q_j^2 ,*

$$\sigma_j = -\frac{q_j^2 - q_j^1}{\log[\kappa_j/p_j(q_j^2) - 1] - \log[\kappa_j/p_j(q_j^1) - 1]}. \quad (47)$$

Proof. The mean-field hiring equation implies

$$\log\left[\frac{\kappa_j}{p_j(q_j)} - 1\right] = \log D_j - \frac{q_j}{\sigma_j}.$$

The slope in productivity rank is $-1/\sigma_j$. Differencing between two ranks yields equation (47). $\square$

With more than two rank bins, linearity of transformed hiring odds and agreement of the intercept with the equilibrium competition index D_j provide overidentifying restrictions.

D.2.3 Worker-side parameters

Define the observable objects

$$G_j(q) = -\alpha_j(q)\beta p_j(q_j) \log[1 - \lambda_j(q)].$$

For each joint cell n , the two occupation equations can be written

$$w_A = b + b_1 \bar{q}_n + \frac{c_A}{p_{nA}} + s_A [\text{logit}(\lambda_{nA}) + G_{nA}] + s_B G_{nB}, \quad (48)$$

$$w_B = b + b_1 \bar{q}_n + \frac{c_B}{p_{nB}} + s_A G_{nA} + s_B [\text{logit}(\lambda_{nB}) + G_{nB}]. \quad (49)$$

Stack the equations as

$$W = X\theta_u, \quad \theta_u = (b, b_1, c_A, c_B, s_A, s_B)'. \quad (50)$$

The A and B rows for cell n are

$$X_{nA} = \left(1, \bar{q}_n, \frac{1}{p_{nA}}, 0, \text{logit}(\lambda_{nA}) + G_{nA}, G_{nB} \right),$$

$$X_{nB} = \left(1, \bar{q}_n, 0, \frac{1}{p_{nB}}, G_{nA}, \text{logit}(\lambda_{nB}) + G_{nB} \right).$$

Proposition 3 (Identification of worker-side parameters). *If X has column rank six, then*

$$\theta_u = (X'X)^{-1}X'W.$$

Thus the conditional application and hiring moments point identify $(b, b_1, c_A, c_B, s_A, s_B)$.

Proof. Equations (48)–(49) give $W = X\theta_u$. Full column rank implies that the null space of X is zero, so no two distinct parameter vectors produce the same moment collection. Premultiplying by $(X'X)^{-1}X'$ gives the inverse. $\square$

Theorem 1 (Mean-field identification). *Under the interiority, rank, observed-primitives, and equilibrium-selection conditions above, the mapping from admissible structural parameters (θ_0, θ) to the collection of aggregate, conditional application, and conditional hiring moments is one-to-one. Its inverse is given sequentially by equations (44)–(45), the directly observed notification and cell-distribution moments, Proposition 2, and Proposition 3.*

The DiD moments are not required for this constructive point-identification result. They overidentify the model and discipline its equilibrium response away from the pre-reform wage.

D.3 GMM moments

Let superscript M denote a model prediction at the pre-reform wage. For each joint cell, the standardised application moments are

$$g_{\lambda_{A,c}}(\theta) = \frac{\text{logit}(\hat{\lambda}_{A,c}) - \text{logit}(\lambda_{A,c}^M(\theta))}{\widehat{\text{se}}[\text{logit}(\hat{\lambda}_{A,c})]}, \quad (50)$$

$$g_{\lambda_{B,c}}(\theta) = \frac{\text{logit}(\hat{\lambda}_{B,c}) - \text{logit}(\lambda_{B,c}^M(\theta))}{\widehat{\text{se}}[\text{logit}(\hat{\lambda}_{B,c})]}. \quad (51)$$

For each distinct occupation-specific rank bin, the hiring moment is

$$g_{p_j,x}(\theta) = \frac{\log(\hat{\kappa}_j/\hat{p}_j(x) - 1) - \log(\hat{\kappa}_j/p_j^M(x;\theta) - 1)}{\widehat{\text{se}}[\log(\hat{\kappa}_j/\hat{p}_j(x) - 1)]}. \quad (52)$$

The estimator also includes pre-reform applications per vacancy:

$$g_{\mathcal{A}_j}(\theta) = \frac{\hat{\mathcal{A}}_j - \mathcal{A}_j^M(\theta)}{\widehat{\text{se}}(\hat{\mathcal{A}}_j)}. \quad (53)$$

Let $\hat{\tau}_{\mathcal{A}}^{DID}$ and $\hat{\tau}_Q^{DID}$ denote the empirical coefficients. The DiD moments are

$$g_{DID,\mathcal{A}}(\theta) = \frac{\hat{\tau}_{\mathcal{A}}^{DID} - \tau_{\mathcal{A}}^M(\theta)}{\widehat{\text{se}}(\hat{\tau}_{\mathcal{A}}^{DID})}, \quad (54)$$

$$g_{DID,Q}(\theta) = \frac{\hat{\tau}_Q^{DID} - \tau_Q^M(\theta)}{\widehat{\text{se}}(\hat{\tau}_Q^{DID})}. \quad (55)$$

Stacking all moments gives $g(\theta)$. The estimator minimizes

$$Q(\theta) = g(\theta)'Wg(\theta).$$

Positive parameters are optimised in logarithms. A constructive estimate based on Propositions 2 and 3 provides starting values; the numerical estimator uses several dispersed starts and retains the solution with the lowest criterion.

D.4 Weighting matrix and specification test

Let $\hat{\Omega}$ be the covariance matrix of the unstandardised empirical moments and let $\hat{S}$ be its diagonal matrix of standard deviations. Because g contains standardised residuals, the

efficient weighting matrix is the inverse correlation matrix

$$W = \left(\hat{S}^{-1} \hat{\Omega} \hat{S}^{-1} \right)^{-1}.$$

We estimate $\hat{\Omega}$ by a common cluster bootstrap that recomputes worker ranks when appropriate, cell probabilities, aggregate moments, and both DiD estimates in every draw. The results in the main text use the resulting inverse correlation matrix.

At the optimum,

$$J = g(\hat{\theta})' W g(\hat{\theta}) \stackrel{a}{\sim} \chi_{m-8}^2,$$

where m is the number of moments. We use its standard overidentification-test interpretation.

D.5 Structural decomposition

For each outcome Y , solve the model at the pre- and post-reform wages:

$$\mathcal{E}^0 = \mathcal{E}(w_A^0, w_B; \hat{\theta}), \quad \mathcal{E}^1 = \mathcal{E}(w_A^1, w_B; \hat{\theta}).$$

These equilibria provide $\lambda_j^r(q)$, $p_j^r(q_j)$, and $u^r(q)$ for $r \in \{0, 1\}$. The occupation-level changes are

$$\Delta Y_j = Y_j(\mathcal{E}^1) - Y_j(\mathcal{E}^0),$$

and the structural DiD is exactly

$$\tau_Y^M = \Delta Y_A - \Delta Y_B.$$

When the empirical regression transforms the outcome, Y_j denotes that transformed outcome before differences are taken.

D.5.1 Channels within occupation B

At each point on the wage path, decompose the application response as

$$\frac{d\lambda_B^u(q)}{dw_A} = -\frac{\lambda_B(q)[1 - \lambda_B(q)]}{s_B} \frac{du(q)}{dw_A}, \quad (56)$$

$$\frac{d\lambda_B^p(q)}{dw_A} = \frac{\lambda_B(q)[1 - \lambda_B(q)]}{s_B} \frac{c_B}{p_B(q_B)^2} \frac{dp_B(q_B)}{dw_A}. \quad (57)$$

For applications per vacancy,

$$\frac{d\mathcal{A}_B^u}{dw_A} = M_B \sum_c \nu_B(c) \frac{d\lambda_B^u(c)}{dw_A}, \quad (58)$$

$$\frac{d\mathcal{A}_B^p}{dw_A} = M_B \sum_c \nu_B(c) \frac{d\lambda_B^p(c)}{dw_A}. \quad (59)$$

If the empirical outcome is $h(\mathcal{A}_B)$, each derivative is multiplied by $h'(\mathcal{A}_B)$.

Average applicant quality is

$$\bar{Q}_B = \frac{\sum_c \nu_B(c) q_B(c) \lambda_B(c)}{\sum_c \nu_B(c) \lambda_B(c)}.$$

For a perturbation $d\lambda_B^k$, where $k \in \{u, p\}$, its directional derivative is

$$d\bar{Q}_B^k = \frac{\sum_c \nu_B(c) (q_B(c) - \bar{Q}_B) d\lambda_B^k(c)}{\sum_c \nu_B(c) \lambda_B(c)}. \quad (60)$$

The signs of these quality channels are unrestricted because they depend on which worker types change their application probabilities.

Integrating along the observed wage change gives

$$\Delta Y_B^k = \int_{w_A^0}^{w_A^1} \frac{dY_B^k(w)}{dw} dw, \quad k \in \{u, p\}. \quad (61)$$

The adding-up restrictions are

$$\Delta Y_B = \Delta Y_B^u + \Delta Y_B^p, \quad \tau_Y^M = \Delta Y_A - \Delta Y_B^u - \Delta Y_B^p.$$

D.6 Numerical solution and inference

For each trial parameter vector, the model is solved at w_A^0 . Worker application probabilities are obtained from the scalar continuation-value problem, conditional application probabilities are integrated using the joint productivity distribution, and hiring probabilities are updated from the mean-field competition indices. The iteration continues until the maximum change in log hiring probabilities is below the numerical tolerance.

The post-reform solution at w_A^1 is initialized from the pre-reform equilibrium. These two equilibria generate the model DiD moments. Numerical accuracy is assessed by tightening the productivity grid, fixed-point tolerance, and relaxation parameter.

For the channel decomposition, we solve the equilibrium on a fine grid between w_A^0 and

w_A^1 . centred finite differences approximate du/dw_A and dp_B/dw_A , which are inserted into equations (56)–(60). We then use trapezoidal integration in equation (61). The difference between the total B response and the sum of the two channels is reported as a numerical adding-up error.

For inference, each bootstrap draw repeats the complete procedure:

1. resample at the clustering level used in the empirical analysis;
2. reconstruct the pre-reform and DiD moments;
3. recalibrate the directly identified parameters;
4. re-estimate the eight behavioural parameters;
5. solve the pre- and post-reform equilibria;
6. compute model fit and the structural decomposition.

Percentile or studentized intervals from this distribution incorporate sampling uncertainty in both the reduced-form and structural stages.