Vacancy Duration and Wages*

Ihsaan Bassier Alan Manning Barbara Petrongolo

January 2025

Abstract

We estimate the elasticity of vacancy duration with respect to posted wages, using data from the near-universe of online job adverts in the United Kingdom. Our research design leverages firm-level wage policies that are plausibly exogenous to hiring difficulties on specific job vacancies, and controls for job and market-level fixed-effects. Wage policies are defined based on external information on pay settlements, or on sharp, internally-defined, firm-level changes. In our preferred specifications, we estimate duration elasticities in the range -3 to -5, which are substantially larger than the few existing estimates.

*Bassier: Centre for Economic Performance, London School of Economics, and Economics, University of Surrey, Guildford GU2 7XH, UK (email: i.bassier@surrey.ac.uk); Manning: Centre for Economic Performance and Department of Economics, London School of Economics, Houghton Street, London WC2A 2AE, UK (email: a.manning@lse.ac.uk); Petrongolo: Centre for Economic Performance, London School of Economics, and Department of Economics, Oxford University, Manor Road, Oxford OX1 3UQ, UK (email: barbara.petrongolo@economics.ox.ac.uk). We thank the Urban Big Data Centre of the University of Glasgow for providing the Adzuna data. Manning and Bassier gratefully acknowledge financial support from the European Research Council (Grant LPIGMANN Number 834455).

1 Introduction

One key concept to evaluate the imbalance of power between workers and firms is the elasticity of labor supply to an individual firm. In a dynamic labor market model, higher wages ease worker recruitment and retention, and the elasticities of these entry and exit margins of labor supply to the firm jointly determine employer market power. There is a large and growing body of work estimating the retention elasticity (see Sokolova and Sorensen, 2021), but evidence on the recruitment margin is more limited. Our paper focuses on the elasticity of vacancy duration to the wage, a measure of the extent to which higher-wage firms can fill job vacancies faster, making recruitment easier. Besides its relationship to market power, the study of the determinants of vacancy duration is interesting in its own right, to complement the vast existing evidence on the process of worker search with corresponding evidence on search duration on the employer side.

This paper uses information from the near-universe of online job adverts in the UK, provided by the Adzuna job-search engine, to estimate the elasticity of vacancy duration to posted wages using a variety of empirical strategies. Our preferred specification leverages within-firm, discrete wage changes, such that a job advert is considerably more attractive just after a wage-change event than just before. Exploiting variation from both externally defined, annual pay settlements and internally defined, sharp wage changes, we estimate a vacancy duration elasticity in the range -3 to -5. The estimates obtained are considerably larger than existing estimates for the vacancy elasticity, as well as OLS estimates in our data, even when controlling for detailed job and firm characteristics. For example, OLS estimates of the vacancy elasticity that control for detailed job titles are in the range -0.2 to -0.4, and estimated elasticities with less rich controls are closer to zero — implying implausibly high markdowns of wages below the marginal revenue product of labor. We argue that a valid identification strategy needs to adequately address potential biases from unobserved job, worker and firm characteristics, measurement error in posted wages, and the possible endogeneity of wage adjustments to perceived hiring difficulties and competitor firms' hiring strategies.

Our work contributes to a small emerging literature on the determinants of vacancy duration. Faberman and Menzio (2018) find that vacancy durations in US survey data are *positively* related to entry wages, and suggest this is driven by omitted worker quality controls. Mueller et al. (2023) link job adverts to matched worker-firm data from Austria, and find that vacancy duration is negatively correlated with a hire's entry wage and with permanent, firm-level wage premia, with estimated elasticities of -0.07 and -0.21, respectively. The latter is very close to the -0.19 estimate that we obtain on a specification conceptually similar to theirs. Our paper is also related to empirical work on directed search, and in particular to Carrillo-Tudela et al. (2023), who test the implications of directed-search models on the determinants of vacancy durations. Finally, our paper is more broadly related to recent work on the wage elasticity of job applications (Azar et al., 2022; Banfi and Villena-Roldan, 2019; Belot et al., 2022), as vacancies that offer higher wages and attract more applicants are expected to fill faster, and on the elasticity of recruitment (Dal Bó et al., 2013; Datta, 2023; Falch, 2017; Hirsch et al., 2022). We contribute to this literature with a novel research design exploiting within-firm variation in wages and vacancy durations on a large, representative set of job adverts.

2 Vacancy duration and employer market power

In a monopsonistic labor market, wages are a markdown on marginal products, where markdowns are inversely related to the wage elasticity of labor supply to an individual firm (see Appendix A for theoretical details). Hence the wage elasticity of firm-level employment is often used to measure employer market power. In a dynamic labor market model, steady-state employment (N) is given by the ratio of recruits (R) to the separation rate (s), i.e. $N = R/s$. The labor supply elasticity to the firm is therefore equal to the difference between recruitment and separation elasticities. The flow of recruits can be expressed as the product of the number of vacancies, V, and the rate at which they are filled, θ, i.e. $R = \theta V$. The probability of filling a vacancy is the inverse of its expected duration d, i.e. $\theta = 1/d$, so:

$$\ln R = \ln V - \ln d.$$

The recruitment elasticity is therefore equal to the difference between the elasticity of the number of vacancies and the duration elasticity:

$$\frac{\partial \ln R}{\partial \ln w} = \frac{\partial \ln V}{\partial \ln w} - \frac{\partial \ln d}{\partial \ln w}. \tag{1}$$

While vacancy posting (V) is typically a choice variable for the firm, the duration of a vacancy is more plausibly driven by worker search responses and is therefore especially informative about the competitiveness of labor markets. Interestingly, Carrillo-Tudela et al. (2023) find that variation in recruitment across firms is predominantly accounted for by the vacancy duration margin rather than variation in vacancy rates. Our analysis focuses on the identification of the duration elasticity with respect to the wage. Research on this margin of the labor supply elasticity to the firm is especially scant, as datasets containing information on offered wages and vacancy durations are rare.

3 Data

3.1 Data sources and cleaning

We use a database assembled by Adzuna, a job search engine that scrapes the universe of job vacancies posted online in the UK from 2017 onwards. The Adzuna data record the stock of vacancies each week, and are used by the Office for National Statistics (ONS) as an indicator of UK economic activity (for details, see ONS, 2023). The total stock of vacancies at the monthly level in Adzuna is on average 1.19 times the vacancy stock in the ONS vacancy survey, and the correlation between the number of vacancies at the industry-quarter level in the two sources is 0.69.[1]

The dataset contains information on the name of the posting firm, location, job title and wage, as well as a free-format job description – i.e. exactly the same information that is visible to jobseekers. Each vacancy has a unique identifier, which allows us to link observations across weeks and measure its duration as the number of weeks it remains

[1]While the ONS vacancy survey could be undercounting vacancies, it is also possible that our dataset picks up duplicate vacancies posted on different websites (despite Adzuna using a de-duplicating algorithm). For robustness, we perform additional de-duplication exercises and obtain very similar results to our main analysis.

posted. Within a vacancy, all characteristics stay constant across weeks. We organize the data into a collection of vacancy spells, defined by the first and last week when a vacancy is posted. We restrict our sample to vacancies first posted during 2017-2019, as later data are affected by the pandemic.

There are about 55 million vacancies advertised during 2017-2019. We exclude vacancies for which information on the date first posted is inconsistent with the number of weeks a vacancy is observed (about 3 million), or with missing information on the posting firm, location, or job title (about 2 million). We observe wages for about two thirds of vacancies, a higher incidence than in many vacancy datasets.[2] Vacancies without wage information have a mean duration of 17 days, which is very similar to the 18-days mean observed in our analysis sample, and a similar occupational composition (see columns 1 and 2 in Table C1). Banfi and Villena-Roldan (2019) estimate a higher elasticity of applications to the wage when remuneration is posted, thus estimates of the duration-wage elasticity in our sample may be somewhat larger than in the universe of vacancies. The wage-posting decision appears to be made overwhelmingly at the job-level (60% of jobs always have a wage posted, and 32% never have a wage posted); given that our empirical specifications include job fixed effects, bias from endogenous posting is not especially concerning.[3] Following these selection criteria, we have a working sample of 32.5 million vacancies. Finally, we extract information about job amenities from the vacancy description field, which may influence duration and also correlate with wages. Pensions, (paid) holidays and bonuses are among most frequent amenities (see Figure C1). Further detail on our data is provided in Appendix B.

We (fuzzy) match the Adzuna dataset with some external sources. First, based on firm names, we match the vacancy data with information of firm-level pay settlements collected by the Labor Research Department's (LRD). We match just over half of LRD agreements, but only 71,000 Adzuna vacancies. Secondly, we match job titles with 4-

[2]For example, Batra et al. (2023) report wage information in 14% of US vacancies, similar to Hazell et al. (2022). Wage information varies considerably by country. Brenčič (2012) finds that the share of adverts posting wages is highest in the UK (86%), followed by the US (18%) and Slovenia (25%). Burning Glass UK data have similar rates of wage posting to Adzuna (e.g. Adams-Prassl et al., 2020).

[3]This would address the finding in Škoda (2022) that the probability of posting a wage is systematically correlated to employer size and occupation. We also replicate all our specifications on the subsample of jobs for which wages are always posted, with very similar results.

digit occupations, obtaining a match for two thirds of observations. Finally, we match firm names into the Orbis database, containing information on industry and other firm characteristics, obtaining a match for about half of the sample.

3.2 Descriptive evidence

We define a job j based on the combination of a given job title, posting firm, and location, defined at the travel to work area (TTWA) level.[4] 70% of vacancies for which industry information can be matched are posted by recruitment or temporary help agencies. For these cases, we do not observe the ultimate employer, and we later investigate systematic differences in duration elasticities between vacancies directly posted by an employer and those posted by such intermediaries. Vacancies are posted relatively evenly across the months of the year (see Figure C2). For most jobs, there are usually only a few vacancies advertised over the sample period: 65% of jobs only feature once, corresponding to 32% of vacancy observations in the sample (see Figure C3). Figure C4 plots the distribution of vacancy duration in weeks. 60% of vacancies are removed within 3 weeks. There is a spike around 4 weeks indicating that firms may tend to leave an advert out for a month. In robustness checks, we estimate a truncated regression which censors vacancy duration at 3 weeks to exploit variation that is least affected by the 4-week bunching, and also include regression controls for the original website where a vacancy was posted. In general, the cost of vacancies incentivises firms to withdraw a listing once a vacancy is filled.

One concern is that some job adverts may be withdrawn by employers without being filled, for example because they give up search or hiring intentions change. This issue has an analogy in the empirical literature on separation elasticities, which often may not distinguish between worker quits and layoffs. In our context, vacancy withdrawal may bias elasticity estimates if the withdrawal incidence is high, and the wage elasticity of filling versus withdrawing vacancies is systematically different. However, studies that have access to information on withdrawals suggest that this bias is unlikely to be large. For example, Van Ours and Ridder (1992) and Mueller et al. (2023) find that 4% and

[4]There are 228 TTWAs in the UK, representing commuting zones.

14% of their respective vacancy samples are withdrawn.

Finally, measuring vacancy duration as the length of time it stays advertised, as opposed to the length of time until the new hire starts work, has the advantage of focusing on the time to find candidates – which should respond to the wage posted via labor market competition – rather than the screening and selection process,[5] and the lag between a job offer and the start of an employment spell (Davis et al., 2014; Van Ours and Ridder, 1992).

Among vacancies with non-missing wage information, 40% post a single wage level, while the rest post a wage range, in which case we use the top of the range as a regressor, but perform robustness tests based on the mid-range value. Most salaries are reported on an annual basis, otherwise we convert them to an annual equivalent if posted hourly or daily. For validation, we compare wages in the Adzuna data to wages in the Annual Survey of Hours and Earnings (ASHE), the UK's most comprehensive source of earnings data. Panel A in figure C5 shows the binned scatterplot of median wages by 3-digit occupation, with an underlying slope coefficient of 0.55. Panel B shows a scatterplot of annual wage changes, with a slope coefficient of 0.71.

4 Baseline estimates

Our first set of estimates is obtained on the full analysis sample (see column 3 in Table C1) by regressing the (log) vacancy duration on the (log) posted wage and a set of controls. Duration of search tends to be longer for more skilled workers, because returns to match quality may be higher for specialized skills (Faberman and Menzio 2018, Amior 2019), and/or high-quality workers are relatively scarce. Inadequate controls for job characteristics may therefore lead to upward-biased elasticity estimates because high-skilled jobs are both harder to fill and pay higher wages. To address this point we introduce job fixed-effects. Second, vacancy duration likely responds to the wage offered by competitor firms, alongside own wage. As own and competitors' wages are possibly affected by correlated shocks, we control for week-by-location (i.e. "market") fixed-effects.

[5] An employer may close a vacancy when they deem to have collected enough applications, and later screen applications to select the successful candidate.

Identification therefore exploits variation in posted wages across subsequent adverts for the same job, net of local wage changes.

Our baseline regression has the form:

$$\ln d_{j,t} = \beta \ln w_{j,t} + \gamma_j + \alpha_{l,t} + \nu_{j,t}, \tag{2}$$

where $d_{j,t}$ denotes the duration of a vacancy for job j at calendar time t, β denotes its elasticity with respect to the posted wage $w_{j,t}$, and γ_j and $\alpha_{l,t}$ denote job and week-by-location (TTWA) fixed-effects, respectively.

Results are shown in Table 1. Column 1 estimates equation (2) on the full sample of vacancies and obtains an elasticity estimate of -0.195. Column 2 trims the wage data to reduce the impact of noise and measurement error by residualizing (log) wages by job and date-by-location fixed-effects and dropping 1% of observations at the extremes of the residuals' distribution. The estimated elasticity grows to -0.369. Figure C6 shows non-parametrically the relationship between residualized wages and duration on the trimmed and used observations, respectively. The negative relationship between wages and duration is much stronger on the trimmed sample, likely due to a higher incidence of measurement error at the extremes of the distribution.[6]

Column 3 includes controls for the wage concept (hourly, daily, annual or unstated), for mentions of non-wage benefits in the salary field, and for job amenities extracted from the free-format job description (see the list shown in Figure C1). None of the added sets of controls makes a difference to the estimated elasticities. This is unsurprising, as we include job fixed-effects, and observable characteristics explain little further variation in salary and duration. Figure C8 considers additional controls, including firm-level characteristics (such as employment growth), all interacted with calendar time, as well as additional specifications based on firm-level wages,[7] first-differences, censored-duration regressions, or the sample of jobs that always post wages. Estimates range between -0.5 and -0.1.

[6]The estimated elasticity tends to grow with the extent of trimming below 10% and shrink thereafter (see Figure C7); we view a 1% trimming as a reasonable baseline choice.

[7]From a log wage regression that additionally controls for location, 4-digit occupation and 4-digit industry (excluding vacancies from worker agencies).

Column 4 in Table 1 shows that, when controlling for firm-by-occupation fixed-effects, the estimated elasticity is considerably smaller than in column 3, which includes job fixed-effects. This importance of controlling for more detailed job characteristics than the usual occupation controls is in line with Marinescu and Wolthoff (2020), who find that the estimated elasticity of job applications to wages switches from negative to positive when occupation controls are replaced with more detailed job title controls.

In summary, our main baseline estimate of -0.37 is significantly and robustly negative, unlike some other estimates reported in the literature (Faberman and Menzio, 2018; Mueller et al., 2023). However, it is much smaller than existing estimates of other margins of the firm-level labor supply elasticity (see, among others, Sokolova and Sorensen, 2021 and Hirsch et al., 2022).[8]

Identification concerns. Our baseline regressions include job fixed-effects, absorbing the role of permanent job characteristics that may be systematically related to duration and wages. However, concerns about reverse causality remain, for example firms may decide to post higher wages on a certain job when they expect it to be harder to fill, leading to an upward bias in the estimated elasticity of duration to wages. Firms may also adjust their recruitment effort or hiring standards with the wage offered.

Various strategies can be used to address this challenge. One possibility consists in relating the duration of a job vacancy to the permanent, firm-level component of wages, which is uncorrelated to idiosyncratic fluctuations in hiring difficulties on a given job. This is the strategy adopted by Mueller et al. (2023), who estimate the elasticity of vacancy duration with respect to the firm-specific component of wages, obtained in an AKM decomposition (Abowd et al., 1999). In their analysis, the use of the AKM firm component, as opposed to a worker's entry wage, raises the estimated elasticity from -0.075 to -0.211. Without matching vacancies to employer-employee data registers, this specification cannot be replicated exactly in our data. However, we can obtain an

[8] As the elasticity of vacancy posting to wages is small (Carrillo-Tudela et al., 2023), if separation and recruitment elasticities are similar (following Manning 2003 and much of the related literature), a recruitment elasticity of 0.37 implies that workers are paid about 40% of their marginal product (based on the markdown formula $w = p\varepsilon/(1+\varepsilon)$, where w and p denote the wage and the marginal revenue product of labor, respectively, and ε denotes the labor supply elasticity, which is equal to the sum of the recruitment and separation elasticities).

estimate of the permanent, firm-level component of wages in the Adzuna data and use this as a regressor in a vacancy duration equation that is conceptually similar to the specification of Mueller et al. (2023). This procedure yields a vacancy elasticity close to theirs, of -0.19 (see Figure C8). Although this specification caters for reverse causality between the wage posted and vacancy duration, it is essentially a firm-level cross-sectional regression, and the estimates obtained may be biased by unobservables correlated with the firm wage effect, e.g. permanent firm-level amenities or systematic worker selection patterns.[9]

Finally, measurement error in the posted wage (a concern also highlighted by Batra et al., 2023 for the US) would attenuate the estimated duration elasticity. One clear indicator of the extent of noise in the wage data is the large impact of 1% trimming on the elasticity estimate in Table 1. Measurement error becomes especially problematic when job fixed-effects are included, as they reduce the signal to noise ratio.

5 Wage-change events

5.1 Research design

Our proposed identification strategy exploits sharp, plausibly exogenous firm-level changes in wages, implying that a firm's job vacancies will be most competitive just after a discrete wage adjustment and least competitive just before, against a backdrop of constant or slowly-evolving amenities and labor markets conditions. To implement this strategy, we leverage the fact that most firms have in place policies to revise wages at regular intervals (in most cases annually), and the associated wage change is typically sizable and applies across all jobs in a firm, hence the timing and magnitude of a wage policy is unlikely to be influenced by contemporaneous shocks to hiring difficulties on a given job.

We follow two complementary approaches, based on what we define as "external" and "internal" measures of wage adjustments. The external measure imports information on pay settlements surveyed by the LRD, an independent, trade-union based, research organisation that collects data on collective agreements. The database contains infor-

[9] See Bassier et al. (2022), who initially use a similar strategy to estimate the separations elasticity.

mation on the firm's name, the dates when wages were adjusted, and the associated, company-wide wage change. Figures C9 and C10 show distributions of agreement dates and magnitudes. To focus on annual wage changes, we restrict to companies where at least 80% of agreements between 2013 and 2019 happen on the first day of the same month every year, and use such wage changes as identifying variation over 2017-2019. We then match LRD and Adzuna data on the company names. In the final regression sample, there are 440 unique firm-level wage adjustments, covering 65,789 vacancies, corresponding to 19,409 jobs across 215 firms (see column 4 in Table C1 for descriptives on this sample). We include control vacancies in this regression sample (as in the procedure of Borusyak et al., 2024), selected as vacancies in firms that do not feature in the LRD database and have no large firm-level wage increases over the full sample period.

The internal measure infers wage-setting events from information on wages posted in the Adzuna sample, by isolating weeks in which there is a discrete wage change, surrounded by weeks without wage changes. We first compute the average wage change for firm f at time t across all advertised jobs j, $\Delta \ln w_{f,t} = \frac{1}{n_f}\sum_j^{n_f}(\Delta \ln w_{j,t})$, where $\Delta \ln w_{j,t}$ denotes the (log) wage difference between the current and the most recent posting of job j (which may have happened many weeks earlier) and $\Delta \ln w_{f,t}$ takes the average of all such changes for each period and firm. Events are defined as any firm-week observations with an average wage increase above 5% and below an implausible 50% ($\Delta \ln w_{f,t} \in [0.05, 0.5]$), and a surrounding 24-week interval without wage increases exceeding 1% ($\Delta \ln w_{f,t+h} < .01$ for $h \in [-12, 11]$ and $h \neq 0$).[10] To limit the influence of any given advert on the definition of a wage event, we restrict this sample to events involving at least three adverts.[11] Nearly 90% of the resulting events have precisely zero wage changes in the surrounding weeks, consistent with the interpretation that these are discrete, firm-level wage changes. Our final regression sample has 283 unique firm-level wage increases, covering 18,856 vacancies, corresponding to 3,461 jobs across 282 firms (see column 6 in Table C1). We also implement a leave-one-out version of the internal measure of wage changes, by relating the duration of a vacancy for job j to the firm-level average wage

[10] The estimates are very robust to alternative thresholds for both treatment and control wage changes.

[11] There is a trade-off between setting a higher number of vacancies per wage event and the reduction in the sample size. We show robustness in Table C5.

change obtained on all adverts at time t, excluding job j. The leave-one-out sample has 1,694 unique wage increases, covering 18,288 vacancies, corresponding to 3,419 jobs across 247 firms. As control firms, we include in the sample firms with a full 24-week span without any wage increase exceeding 1% ($\Delta \ln w_{f,t+h} < .01$ for $h \in [-12, 11]$).

Although the external and internal definitions of wage changes are based on a common idea to identify firm-level wage policies, the two approaches have different samples and strengths and weaknesses. The internal measure of wage events has, partly by construction, a stronger first-stage effect on job-level wages, which improves statistical power and allows us to identify the duration elasticity in an event-study framework. On the other hand, the external measure is more likely to single out firm-level wage policies. For example, Figure C9 shows that pay settlements are concentrated on certain dates, such as 1 April or 1 January, while the internally-defined settlements are spread out across the year.

Relative to the baseline sample, the LRD-matched vacancies cover a higher proportion of low-skill occupations, with lower wages and shorter durations (see columns 3 and 4 in Table C1). These differences reflect the over-representation of collective agreements in LRD-matched vacancies. To some extent, this pattern is also found in the sample of internally-defined wage events (column 6). Control vacancies for each sample (columns 5 and 7, respectively) are more similar to those in the full sample. We address concerns about differential trends in the duration of treated and control vacancies with alternative strategies. First, for the sample of internally-defined wage events, we estimate dynamic effects in an event-study design. The results support the hypothesis of parallel trends. Second, for robustness, we show estimates on treated-only samples, purely exploiting variation from the magnitude and timing of a wage event. Finally, we use a matched sample of control firms based on treated firm covariates.

5.2 Estimates based on external information on pay settlements

Table 2 presents estimates on a sample that includes firms covered in the LRD database and the corresponding control firms. We first show OLS specifications in this reduced sample, controlling for job and location-specific time trends.[12] We obtain an estimate of about -0.1 on the raw wage data (column 1), growing to about -0.4 when we introduce 1% trimming on the wage residuals (column 2). These estimates are close to the corresponding estimates in columns 1 and 2 of Table 1, hence very similar specifications on the different samples yield similar results.

Columns 3-5 show results from IV estimates that use external wage agreement as an instrument for the wage posted, $w_{j,t}$. Column 3 shows a first stage estimate close to 0.5. As expected, this is below 1, reflecting that the the firm's headline pay increase typically does not cover all workers in the same way. The reduced-form estimate is about -2.2, with a resulting IV estimate of about -4.8, and a first-stage F-stat above 50. We report both conventional standard errors[13] and the conservative but more robust Anderson-Rubin confidence intervals for the IV estimates, to cater for potentially weak instruments (Andrews et al., 2019; Lee et al., 2022). This elasticity estimate is about one order of magnitude larger than the baseline estimate of column 2, and highlights the importance of exploiting largely exogenous wage events to avoid biases in OLS estimates.

While variation in measured amenities coincides with external wage events for 8% of jobs, column 4 shows that controlling for posted amenities does not affect the wage elasticity.[14] Figure C11 shows robustness on the extent of trimming, reporting very similar IV estimates except for extremely high levels of trimming, for which confidence intervals are much larger.

Another potential concern is that, while the timing of wage settlements is largely predetermined in this sample, the magnitude of the wage change may be endogenous to a firm's current hiring performance. To address this, column 5 in Table 2 only exploits the timing of wage increases – not the magnitude – using a step-wise dummy as a wage instrument. The estimates imply that firms who sign a pay settlement on average raise their wages by 1.4% (first stage), and see a reduction in vacancy duration by 8.7% (reduced form). The corresponding IV estimate is about -6.

[12]To increase statistical power, we include detailed location-specific quadratic trends for calendar weeks, rather than unrestricted TTWA-by-week fixed effects, as in Table 1. See Table C3 for robustness.

[13]Angrist and Kolesár (2023) argue that distortions in standard errors are small unless endogeneity is "extraordinarily high".

[14]Note however that we do not observe recruitment effort, so we cannot directly control for this margin.

5.3 Estimates based on internally-defined wage-change events

Table 3 shows elasticity estimates based on internally-defined wage changes. All regressions include fixed effects for time-by-location as well as "events", i.e. the 24-week spell that surrounds a wage change. We start by showing estimates with baseline controls in columns 1 and 2 and obtain slightly smaller elasticity estimates than in columns 1 and 2 of Table 1. We next instrument the wage posted on a given job advert at time t, $\ln w_{j,t}$, with the firm-level wage change, $\Delta \ln w_{f,t}$: this is a step function equal to 0 before the event and equal to the firm-wide wage change afterwards. Column 3 shows a relatively high and precise first stage estimate of 0.8, a reduced-form effect of -2.7, and an IV elasticity estimate of about -3.3. Column 4 controls for amenities on the adverts, and column 5 uses the leave-one-out version of the instrument: $\Delta \ln w_{f,t} = \frac{1}{n_f - 1} \sum_{k \neq j}^{n_f} (\Delta \ln w_{k,t})$. As expected, the first-stage estimate is lower, at about 0.6, as is the F-statistic. However, the resulting IV estimate is still negative and highly significant at -4.3, with negative bounds.

Overall, we find that duration elasticities are considerably larger when using cleaner research designs, with estimates in the range of -3 to -5. This implies a markdown on wages below marginal product in the range 15-25%, which is in line with markdowns estimated from separations elasticities.

We think our IV estimates are much larger than the baseline estimate in section 4 primarily because of measurement error in posted wages.[15] Table C2 provides some support for this, based on first-differenced estimates on the external IV sample. The OLS estimate is -0.7 (column 1), much lower than the IV estimate of -5.7 in column 3. Column 2 uses as IV an indicator for wage growth above 1%: this isolates larger wage changes, which may be less prone to measurement error. The estimate is about 3-4 larger than the OLS estimate, and half of the preferred IV estimate, suggesting that a major source of bias in OLS is measurement error.[16] The other likely source of bias, endogenous

[15] Sokolova and Sorensen (2021) detect an analogous difference in estimates for separations elasticities, in the range -0.5 to -1.5 on cross-sectional specifications, and in the range -2 to -5 with better-identified research designs.

[16] Some indirect evidence on the role of measurement error is provided by the relevance of trimming in our estimates (see Table 1). Assuming that trimming primarily reduces mismeasurement in the wage, and setting the "maximum" R-squared conservatively at 0.01 based on the calibration exercises in Mueller

wage adjustment, perhaps explains the remaining difference.

5.4 Event-study estimates

We next consider dynamic effects of discrete wage changes in an event-study design. Based on the internal wage-change measure used in Section 5.3, the following event-study specification relates the duration of vacancies advertised in the 24 weeks around a wage-change event to the magnitude of the wage change (Cengiz et al., 2019; de Chaisemartin et al., 2022):

$$\ln d_{j,t+\tau} = \sum_{u=-12, u\neq -1}^{u=11} \beta_u \Delta \ln w_{f,t} \times 1\{\tau = u\} + \gamma_j + \alpha_{l,t} + \psi_\tau + \nu_{j,t+\tau}, \tag{3}$$

where $t = 0$ denotes the time of the event, $\tau \in [-12, 11]$ denotes the 24-week interval around it, and ψ_τ denotes event-time fixed effects. Equation (3) leverages wage-change events at the firm level ($\Delta \ln w_{f,t}$) and represents the dynamic equivalent of the reduced-form specifications shown in Table 3. Standard errors are clustered at the firm level. Figure C12 in the Appendix shows average wage changes on job adverts around an event, with few wage changes happening before, a sharp jump in wages on the event date, and stable wages thereafter.

Figure 1 plots estimates of the event dummies β_u, which are interpreted as the duration elasticity of a vacancy posted at time $t+u$, with respect to the wage change that took place at time t. The estimates are close to zero before the event (only one of them is significantly different from zero), supporting the parallel-trends assumption and suggesting that these large wage increases are not systematically related to prior dynamics in vacancy durations. The average post-period coefficient is negative and significant, whose average corresponds to the reduced-form estimate reported in column 3 of Table 3. The dynamic pattern is noisy, but is suggestive of some initial delay in the response of vacancy duration to wage changes.

For robustness, we show in Figure C13 estimates that additionally control for week-by-treatment fixed effects: these exclusively leverage variation in the magnitude of firm-

et al. (2023), the Oster (2019) approach yields an extrapolated unbiased durations elasticity of -6.95, which is closer to our IV estimate.

level wage increases. As the event-study sample is highly unbalanced, estimates shown in Figure C14 use alternative minima for the number of vacancies used to build wage-change events. Finally, Figure C15 shows estimates that use leave-one-out firm-level wage changes.

5.5 Robustness and heterogeneity in duration elasticities

We perform extensive robustness analyses. Table C3 considers additional time-varying labor market controls – including local occupation controls – to capture the role of hiring difficulties and/or competitor firms' hiring strategies. The elasticities based on the external IV (columns 1-6) are in the range -3 to -5, similar to the main estimates of Table (2), although we are underpowered in the specification of column 1 that flexibly controls for location-by-week FE (see also footnote 12). The specifications based on the internal IV (columns 7-8) also provide similar estimates to the main specifications of Table (3)

We consider next alternative control samples in Table C4. First, we restrict the regression sample to treated observations. For the external IV especially, this addresses concerns about missing treatments in the LRD database. We then include matched controls (using wages, occupation and location-specific trends as covariates) and report estimates based on propensity-score weights and nearest neighbour matching (Goldschmidt and Schmieder, 2017; Roth et al., 2023). The estimates are similar to those obtained on the main specification for both the external and internal IVs; estimates are also similar if reweighted to match the more representative baseline sample.

The IV estimates are robust to including fixed effects for the posting website (columns 1 and 4 in Table C5), and to using mid-range wages as regressors whenever adverts report a wage range (columns 2 and 5). The external IV allows us to check for bias from endogenous wage-posting by including matched firms which did not post wages, and the reduced form estimate is very similar (column 3). The internal IV is also robust to defining events based on wage changes from at least ten vacancies (as opposed to three, see columns 5 and 6).

Finally, Table C6 shows that estimates are robust to dropping potential duplicate adverts in the Adzuna database: we remove "duplicate" vacancies sharing the same title, location and week in columns 1-2 and 5-6 (e.g. if an employer posts anonymously on different websites, including via recruitment agencies), and we combine durations from vacancies for the same job and with the same firm that close and reopen within 7 days in columns 3-4 and 7-8 (implicitly assuming that the first vacancy was not filled).

We finally investigate heterogeneity in duration elasticities using the internally-defined measure of wage changes.[17] Figure 2 first shows that the magnitude of the duration-wage elasticity is significantly higher for areas with above-median vacancy rates (-3.8) than in areas with below-median rates (-0.9), suggesting that slacker markets are less competitive as workers have fewer outside options. Second, we find that duration elasticities are larger for vacancies posted by agencies (-4.1) than for those posted by direct employers (-1.6), though the difference is not statistically significant. Possibly, workers find it easier to compare wages on agencies' websites, which collate adverts from several employers in more standardized formats, making behavior more sensitive to wage differences. Third, estimates by skill level groups do not show significant differences.

6 Conclusions

This paper has presented evidence that firms that pay higher wages find it easier to fill vacancies. Our preferred specifications, leveraging variation from firm-level wage policies that are plausibly exogenous to hiring difficulties on specific job vacancies, deliver elasticities for vacancy duration to wages in the range -3 to -5. These estimates are in the same ballpark as well-identified estimates of the separations elasticity from existing studies, and are much bigger than prior estimates of the vacancy elasticity.

[17] The sample based on the external IV does not allow enough power to investigate heterogeneous effects.

Table 1: Baseline estimates of the duration-wage elasticity

	(1)	(2)	(3)	(4)
Log wage	-0.195***	-0.369***	-0.383***	-0.178***
	(0.001)	(0.002)	(0.002)	(0.001)
Trimmed		✓	✓	✓
Controls				
Date × TTWA FE	✓	✓	✓	✓
Job FE	✓	✓	✓	
Wage concept and amenities			✓	
Firm × Occupation FE				✓
No. vacancies (M)	21.62	21.24	21.24	19.91
No. jobs (M)	5.98	5.94	5.94	5.53

Notes. The Table shows results from regressions of log vacancy duration on log wages and the indicated controls (see equation 2). The sample includes cleaned vacancy data for 2017–2019, restricted to jobs with at least 2 adverts. All specifications include fixed effects for date by travel-to-work area (TTWA). Columns 2-4 exclude observations in the 1% tails of residualized wages. Controls in column 3 include dummies for the wage concept (hourly, daily, annual), and for non-wage benefits and amenities as extracted from advert descriptions. Occupation fixed-effects in column 4 are at the 4-digit level. Standard errors are reported in brackets. The number of observations and number of vacancies are shown in millions.

Table 2: Estimates of duration-wage elasticity based on external pay settlements

	(1)	(2)	(3)	(4)	(5)
First stage			0.465***	0.464***	0.014***
			(0.064)	(0.064)	(0.002)
Reduced form			-2.238**	-2.242**	-0.087***
			(0.932)	(0.932)	(0.027)
Main equation	-0.110***	-0.425***	-4.816**	-4.827***	-6.046***
	(0.020)	(0.062)	(2.146)	(2.147)	(2.047)
A-R CI			[-9.32,-0.82]	[-9.33,-0.83]	[-10.30,-2.23]
F-stat			53.149	53.142	63.039
Trimmed		✓	✓	✓	✓
Pay set. IV			✓	✓	✓
Timing only					✓
Controls					
Job FE	✓	✓	✓	✓	✓
Location trends	✓	✓	✓	✓	✓
Amenities				✓	✓
Vacancies	392773	389167	389167	389167	389167
Jobs	130972	130297	130297	130297	130297

Notes. The Table shows results from regressions of log vacancy duration on log wages and the indicated controls. The sample includes cleaned vacancy data for 2017–2019, restricted to firms that can be matched to wage agreements in the LRD database and corresponding control firms. All specifications include a quadratic trend in calendar week interacted with TTWA fixed effects. Columns 1 and 2 show baseline OLS estimates on this sample. Column 3 instruments the wage posted by the firm-level pay settlement from the LRD database. Column 4 additionally controls for amenities, measured as categories of benefits listed in the job advert (see Appendix B). In column 5, the instrument is a dummy variable for a pay settlement. Columns 2-5 exclude observations in the 1% tails of non-zero, residualized wage changes. A-R CI indicates the Anderson-Rubin confidence interval for the IV estimate, where a missing bound indicates an unbounded interval on that side. Standard errors are reported in brackets.

Table 3: Estimates of duration-wage elasticity, based on internally-defined wage-change events

	(1)	(2)	(3)	(4)	(5)
First stage			0.824***	0.824***	0.568***
			(0.057)	(0.057)	(0.125)
Reduced form			-2.675***	-2.669***	-2.415***
			(0.793)	(0.790)	(0.748)
Main equation	-0.074***	-0.124***	-3.247***	-3.239***	-4.251***
	(0.004)	(0.005)	(1.000)	(0.995)	(1.419)
A-R CI			[-6.13,-0.67]	[-6.11,-0.68]	[.,-1.04]
F-stat			209.010	208.605	20.652
Trimmed		✓	✓	✓	✓
Firm wage IV			✓	✓	✓
Leave-one-out					✓
Controls					
Date × TTWA FE	✓	✓	✓	✓	✓
Job FE	✓	✓			
Event FE			✓	✓	✓
Amenities				✓	✓
Vacancies	236592	232226	227704	227704	210032
Jobs	12156	12143	11163	11163	10232

Notes. The table shows coefficients from regressions of log vacancy duration on log wage. The sample includes cleaned vacancy data for 2017–2019, restricted to firms experiencing an internally-defined wage-change event and corresponding control firms. All specifications include fixed effects for date by TTWA. Columns 1 and 2 show baseline OLS estimates on this sample. Columns 3-4 instrument the wage posted by the mean and leave-one-out mean firm-level wage change, respectively. Column 4 additionally controls for amenities, measured as categories of benefits listed in the job advert (see Appendix B). Event FE refer to weeks in each 24-week window around each event. Columns 2-5 exclude observations in the 1% tails of non-zero, residualized wage changes. A-R CI indicates the Anderson-Rubin confidence interval for the IV estimate, where a missing bound indicates an unbounded interval on that side.

Figure 1: Event-study estimates of vacancy elasticities

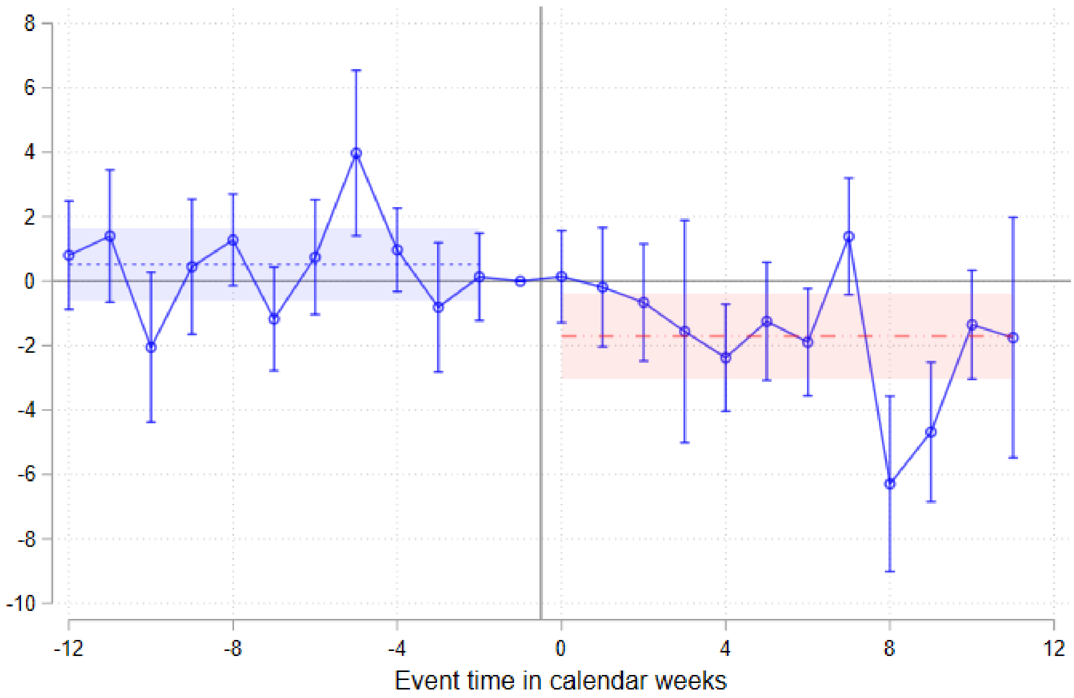

Notes. The Figure shows weekly coefficients from an event-study regression of log vacancy duration on log firm-level wage changes, including fixed effects for jobs, event-time, and date by TTWA (see equation 3). The sample includes cleaned vacancy data for 2017–2019, restricted to firms experiencing an internally-defined wage-change event and corresponding control firms, and excluding observations in the 1% tails of non-zero, residualized wage changes. Event-time zero refers to the firm-level wage change. The horizontal dashed lines show the averaged effects for the pre- and post-event periods. Vertical bars and shaded areas represent 95% confidence intervals.

Figure 2: Heterogeneity analysis on duration elasticities

Notes. The figure shows coefficient magnitudes from regressions of log vacancy duration on log wage, instrumented by internally-defined wage-change events, controlling for job and date by TTWA fixed effects. All coefficients are negative. The sample includes cleaned vacancy data for 2017–2019, restricted to firms experiencing an internally-defined wage-change event and corresponding control firms, and excluding observations in the 1% tails of non-zero, residualized wage changes. Separate regressions are estimated for heterogeneity dimension, with interactions for the subsample indicated for each bar. 95% confidence intervals refer to the difference with respect to the base category (left-most bar of each group). The baseline sample has 11,140 jobs (matching column 3 of table 3). Vacancy density is measured as the number of posted vacancies over regional employment. Sample sizes: 998 jobs below median; 8,792 jobs above median. Worker agency status is determined by the industry code merged from ORBIS data, and indicates recruitment or temporary worker agencies. Samples size: 3,776 jobs for agencies and 1,614 jobs for non-agencies. Occupations are drawn from 1-digit SOC2020 codes. *High skill* indicates managers, professionals and associate professionals (2,564 jobs); *mid-skill* indicates administrative, trade and service occupations (2,073 jobs); and *low skill* indicates sales, operators and elementary occupations (2,657 jobs).

References

Abowd, John M., Francis Kramarz, and David N. Margolis, "High Wage Workers and High Wage Firms," *Econometrica*, 1999, *67*, 251–333.

Adams-Prassl, Abi, Maria Balgova, and Matthias Qian, "Flexible work arrangements in low wage jobs: Evidence from job vacancy data," Discussion Paper DP15263, CEPR 2020.

Amior, Michael, "Education and Geographical Mobility: The Role of the Job Surplus," Discussion Paper 1616, Centre for Economic Performance 2019.

Andrews, Isaiah, James H Stock, and Liyang Sun, "Weak instruments in instrumental variables regression: Theory and practice," *Annual Review of Economics*, 2019, *11*, 727–753.

Angrist, Joshua and Michal Kolesár, "One instrument to rule them all: The bias and coverage of just-id iv," *Journal of Econometrics*, 2023.

Azar, José, Ioana Marinescu, and Marshall Steinbaum, "Labor market concentration," *Journal of Human Resources*, 2022, *57* (S), S167–S199.

Banfi, Stefano and Benjamin Villena-Roldan, "Do high-wage jobs attract more applicants? Directed search evidence from the online labor market," *Journal of Labor Economics*, 2019, *37* (3), 715–746.

Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu, "Monopsony in Movers The Elasticity of Labor Supply to Firm Wage Policies," *Journal of Human Resources*, 2022, *57* (S), S50–s86.

Batra, Honey, Amanda Michaud, and Simon Mongey, "Online Job Posts Contain Very Little Wage Information," Discussion Paper, National Bureau of Economic Research 2023.

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

Bó, Ernesto Dal, Frederico Finan, and Martín A Rossi, "Strengthening state capabilities: The role of financial incentives in the call to public service," *The Quarterly Journal of Economics*, 2013, *128* (3), 1169–1218.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *The Review of Economic Studies*, 2024, *91* (6), 3253–3285.

Brenčič, Vera, "Wage posting: evidence from job ads," *Canadian Journal of Economics/Revue canadienne d'économique*, 2012, *45* (4), 1529–1559.

Carrillo-Tudela, Carlos, Hermann Gartner, and Leo Kaas, "Recruitment policies, job-filling rates and matching efficiency," *Journal of the European Economic Association*, 2023, pp. 2413–2459.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The effect of minimum wages on low-wage jobs," *The Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Datta, Nikhil, "The measure of monopsony: the labour supply elasticity to the firm and its constituents," Discussion Paper 1930, Centre for Economic Performance 2023.

Davis, Steven J, Christof Röttger, Anja Warning, and Enzo Weber, "Job Recruitment and Vacancy Durations in Germany," Discussion Paper 481, University of Regensburg 2014.

de Chaisemartin, Clément, Xavier D'Haultfoeuille, Félix Pasquier, and Gonzalo Vazquez-Bare, "Difference-in-Differences Estimators for Treatments Continuously Distributed at Every Period," Discussion Paper 2201.06898, arXiv 2022.

Faberman, Jason and Guido Menzio, "Evidence on the Relationship between Recruiting and the Starting Wage," *Labour Economics*, 2018, *50*, 67–79.

Falch, Torberg, "Wages and recruitment: Evidence from external wage changes," *ILR Review*, 2017, *70* (2), 483–518.

Goldschmidt, Deborah and Johannes F Schmieder, "The rise of domestic outsourcing and the evolution of the German wage structure," *The Quarterly Journal of Economics*, 2017, *132* (3), 1165–1217.

Hazell, Jonathon, Christina Patterson, Heather Sarsons, and Bledi Taska, "National wage setting," Discussion Paper 30623, National Bureau of Economic Research 2022.

Hirsch, Boris, Elke J Jahn, Alan Manning, and Michael Oberfichtner, "The wage elasticity of recruitment," Discussion Paper 1883, Centre for Economic Performance 2022.

Lee, David S, Justin McCrary, Marcelo J Moreira, and Jack Porter, "Valid t-ratio Inference for IV," *American Economic Review*, 2022, *112* (10), 3260–90.

Manning, Alan, *Monopsony in motion: Imperfect competition in labor markets*, Princeton: Princeton University Press, 2003.

Marinescu, Ioana and Ronald Wolthoff, "Opening the black box of the matching function: The power of words," *Journal of Labor Economics*, 2020, *38* (2), 535–568.

Mueller, Andreas, Damian Osterwalder, Josef Zweimüller, and Andreas Kettemann, "Vacancy Durations and Entry Wages: Evidence from Linked Vacancy-Employer-Employee Data," *The Review of Economic Studies*, 2023, *3*, 1807–1841.

ONS, "Using Adzuna data to derive an indicator of weekly vacancies: Experimental Statistics," https://tinyurl.com/pft5hkzm 2023. Accessed: 2023-07-07.

Oster, Emily, "Unobservable selection and coefficient stability: Theory and evidence," *Journal of Business & Economic Statistics*, 2019, *37* (2), 187–204.

Ours, Jan Van and Geert Ridder, "Vacancies and the recruitment of new employees," *Journal of Labor Economics*, 1992, *10* (2), 138–155.

Roth, Jonathan, Pedro HC Sant'Anna, Alyssa Bilinski, and John Poe, "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 2023.

Škoda, Samuel, "Directing job search in practice: mandating pay information in job ads," Working Paper, University of Zurich, Zurich 2022.

Sokolova, Anna and Todd Sorensen, "Monopsony in labor markets: A meta-analysis," *ILR Review*, 2021, *74* (1), 27–55.

A Theoretical Appendix

We outline here the relationship between the vacancy duration elasticity, the labor supply elasticity, and the wage markdown. Consider a stylized version of a dynamic monopsony model (based on Manning, 2003, section 2.2). Denote by $R(w_t)$ the flow of recruits in a firm at time t, where w_t denotes the wage paid, and by $s(w_t)$ the separation rate. Employment at time t, denoted N_t, then follows the dynamic process:

$$N_t = (1 - s(w_t))N_{t-1} + R(w_t). \tag{4}$$

Equation (4) may be written in steady state as $N = (1 - s)N + R$, thus $N = R/s$, and the long-run labor supply elasticity is equal to the difference between recruitment and separation elasticities:

$$\varepsilon^{LR} = \varepsilon^{rec} - \varepsilon^{sep}. \tag{5}$$

The flow of recruits to the firm can be written as $R = V/d$, where V is the number of vacancies and d is their average duration. Therefore the recruitment elasticity can be written as:

$$\varepsilon^{rec} = \varepsilon^{V} - \varepsilon^{d}. \tag{6}$$

Our paper estimates ε^d, the elasticity of vacancy duration with respect to the wage. Knowledge of ε^V is therefore needed to work out the recruitment elasticity ε^{rec}. Carrillo-Tudela et al. (2023) find that variation in recruitment across firms is predominantly accounted for by the vacancy duration margin rather than variation in vacancy rates. If the elasticity of the vacancy rate is zero, the duration elasticity equals the recruitment elasticity.

The above links the duration elasticity to the recruitment and long-run labor supply elasticities; we next show how this can be linked to the markdown of wages on marginal product by considering the firm's optimal wage policy.

At any date t the state variable for the firm is last period's workforce N_{t-1}. Define a value function $\Pi(N_{t-1})$ as the maximized discounted value of future profits from date t onwards. Using dynamic programming arguments, we can write:

$$\Pi(N_{t-1}) = \max_{w_t, N_t} Y(N_t) - w_t N_t + \delta \Pi(N_t), \tag{7}$$

where δ is the discount factor. This needs to be maximized subject to the dynamic labor supply curve (4). Taking the derivative of (7) with respect to w_t and factoring in the dependence of N_t on w_t leads to the following first-order condition:

$$[Y'(N_t) - w_t + \delta \Pi'(N_t)] \frac{\partial N_t}{\partial w_t} - N_t = 0. \tag{8}$$

There is also the envelope condition, which allows us to obtain the following derivative of the value function (7):

$$\Pi'(N_{t-1}) = [Y'(N_t) - w_t + \delta \Pi'(N_t)] \frac{\partial N_t}{\partial N_{t-1}}. \tag{9}$$

In steady-state, where wages and employment are constant, we can solve (9) for $\Pi'(N)$:

$$\Pi'(N) = \frac{\frac{\partial N_t}{\partial N_{t-1}}[Y'(N) - w](1-s)}{1 - \delta \frac{\partial N_t}{\partial N_{t-1}}}. \tag{10}$$

Substituting this into (8) and using $\frac{\partial N_t}{\partial N_{t-1}} = 1 - s$ gives the markdown equation:

$$\frac{Y'(N) - w}{w} = \frac{1 - \delta(1-s)}{s\varepsilon^{LR}}. \tag{11}$$

Equation (11) implies that says that the gap between the marginal product and the wage decreases with ε^{LR}, which in turn depends on the duration elasticity (our objective of interest).

B Data Appendix

B.1 Details Additional information on data construction

Further notes on data cleaning

- *Durations:* For vacancies posted in late 2019 we use 2020 data to measure their completed duration. Virtually all vacancy durations are shorter than 2 months, thus this procedure does not extend our coverage to the pandemic period.
- *Wages:* Within the cleaned sample, the wage range is relatively narrow. For example, the minimum is less than half of the maximum in only a tiny proportion (2%), and it is within 30% of the maximum for 90% of cases. We exclude a tiny number of vacancies with implausibly low (below £10,000) or high (above £1bn) annual salaries. The wage is posted as hourly (21% of vacancies), daily (8%), annual (53%) or unstated, and occasionally mentions non-wage benefits in the salary field (13%).
- The Adzuna records do not contain information for December 2019.

Matching on external data sources In all data sources, we standardize the company names as much as possible using the same set of rules, i.e. by setting all letters to uppercase, standardizing abbreviations (such as "LTD" or "INC"), and removing punctuation as well as common words such as "the". We then merge firm names across the data sources based on the degree of match between any two firm names in the dataset. We accept perfect matches, then filter based on the "soundex" code of each word, i.e. sounded out components of the firm names, and finally manually match for a small proportion.

Firm pay settlements. We use the annual pay settlement reports from the Labour Research Department from 2014 to 2020, and read the text into a dataset by extracting fields for company, date of pay settlement, percentage wage raise, and the full description. As mentioned in text, over half of the agreements are successfully merged into the vacancy data. Most vacancy data are not matched because they do not have corresponding wage agreements, and some agreements are not matched because firms do not always post vacancies. The percentage wage raise within each settlement exhibits some variation (e.g. by experience), and since we cannot match on all of these we use only the headline increase. For example, the 2017 agreement from one of the UK's largest retailers, Poundland, gave a minimum of 3% increase and average 4.2% (headline), with slightly higher for young sales assistants. The use of the headline, firm-level settlement as an IV is an intention-to-treat estimate, as noted in Section 5.2 of the paper.

Orbis database. We download information on all UK companies with at least 20 employees in the Orbis database, containing fields such as industry, sales turnover, profit, employment, and cost of employment. We use these as controls in the regression analysis. As mentioned in text, about half of firms in the vacancy data obtain a match in the Orbis database.

Job amenities We extract amenities from the free-text description field of each vacancy posting using the "tidytext" package in R. We begin by separating the free-text field into its component words, then exclude extremely common words (e.g. "the" or "a"), then reduce each word to its root, and finally merge this to a list of predefined amenities. This pre-defined list is done on an ad hoc basis, first by querying ChatGPT for a list of workplace amenities, and then supplementing this with additions from a manual search of 200 vacancies.

The most common amenities are shown in figure C1. Amenities are grouped into six categories which are used as controls in the analysis, namely career, financial, recreation, health, work-life balance, and workplace. Since signal may vary (e.g. training may be *required* rather than *provided*), as controls in regressions we check robustness to different combinations of these including broader controls and narrower versions which focus on amenities with stronger signal (such as *pension* or *bonus*).

Practically, controlling for amenities makes little difference because just 16% have any variation in amenities within jobs. For example, only 4% of jobs in the baseline sample have any variation in pension provision, and similarly for bonus. Within jobs, amenities are significantly related to the wage (positive for bonus and health-related, negative for pension, recreation, career or workplace-related), and, conditional on the wage, *positively* related to duration. In the event study analysis, there is co-variation in amenities for 10% of cases, i.e. comparing the vacancy at event-date 0 to the previous vacancy. The variation primarily comes from amenities related to the workplace (increase) and recreation (decrease), with only 1-2% related to financials (e.g. pension or bonus), career or health.

C Additional Figures and Tables

Figure C1: Amenities mentioned in vacancy postings

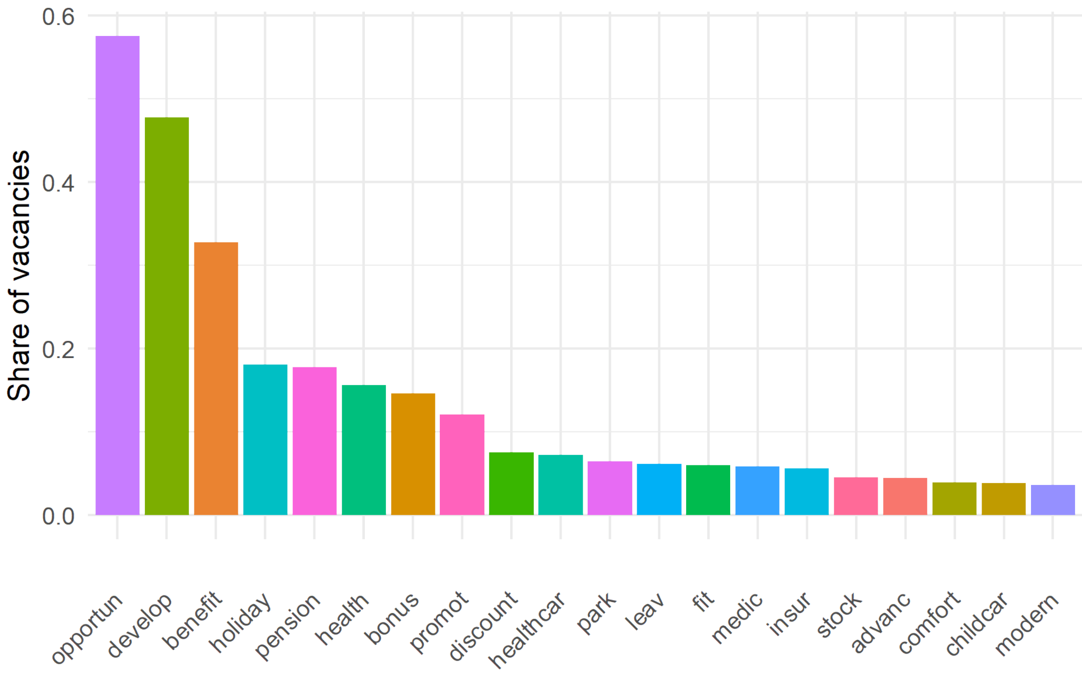

Notes. The figure shows the share of vacancies that mention each amenity "stem" word in the free-format description, e.g.: promot (promote, promotion), medic (medication, medical), insur (insure, insurance), advanc (advance, advancement), etc. The sample includes cleaned Adzuna vacancies for 2017–2019 (see column 2 in Table C1).

Figure C2: Distribution of vacancy postings over time

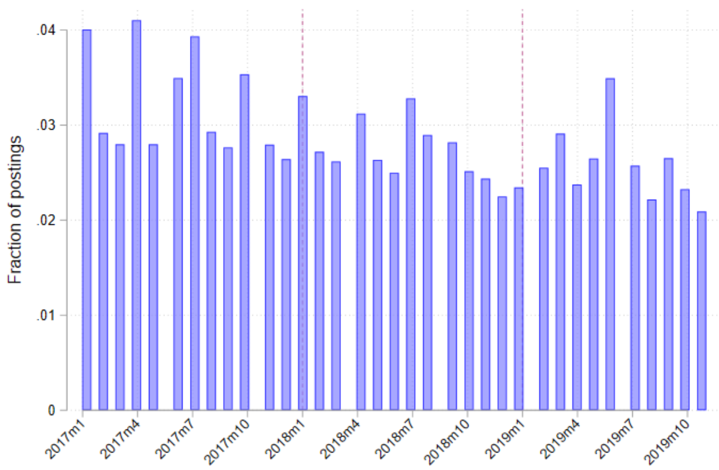

Notes. The figure shows the fraction of vacancy postings by week. The sample includes cleaned Adzuna vacancies for 2017–2019 (see column 2 in Table C1).

Figure C3: Distribution of the number of vacancies per job

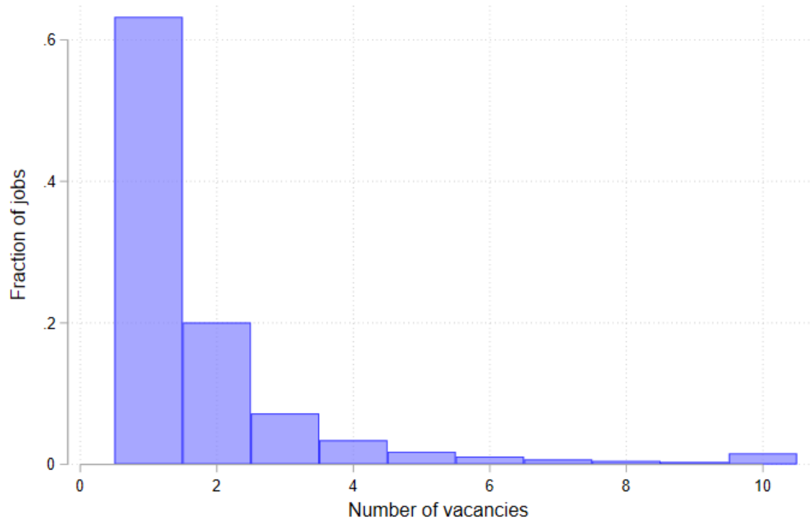

Notes. The figure shows the fraction of jobs by number of vacancies posted over the sample period. A job is defined by the combination of a job title, firm name, and location (TTWA). The right-most bar refers to jobs with 10 or more vacancies. The sample includes cleaned Adzuna vacancies for 2017–2019 (see column 2 in Table C1).

Figure C4: Distribution of vacancy duration

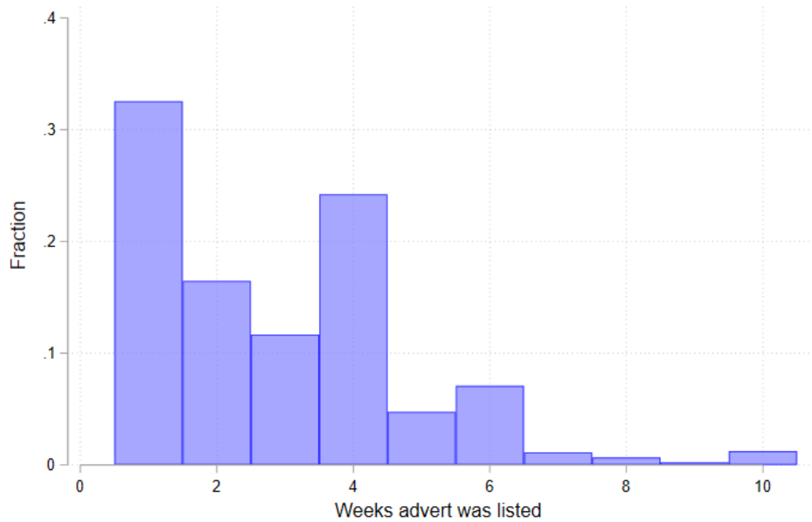

Notes. The figure shows the fraction of vacancies by the number of weeks they are posted. The right-most bar refers to durations of 10 weeks or longer. The sample includes cleaned Adzuna vacancies for 2017–2019 (see column 2 in Table C1).

Figure C5: Wages in Adzuna and ASHE data

(a) **Levels**

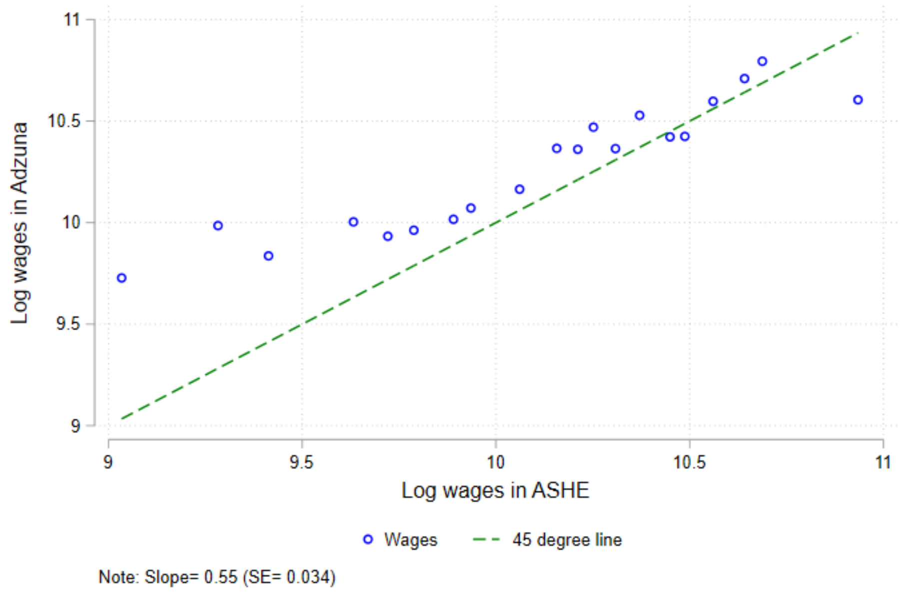

(b) **Differences**

Notes. The figure shows a binned scatter plot of occupation-level wages (3-digit) in Adzuna and ASHE data. Panel (a) compares log wages across the two datasets and panel (b) plots the within-occupation annual wage change. The sample includes cleaned Adzuna vacancies for 2017–2019 (see column 2 in Table C1). The ASHE is an employer-based survey, covering a 1% random sample of employees in the UK. Data are weighted by the number of workers in each cell.

Figure C6: Vacancy durations and wages

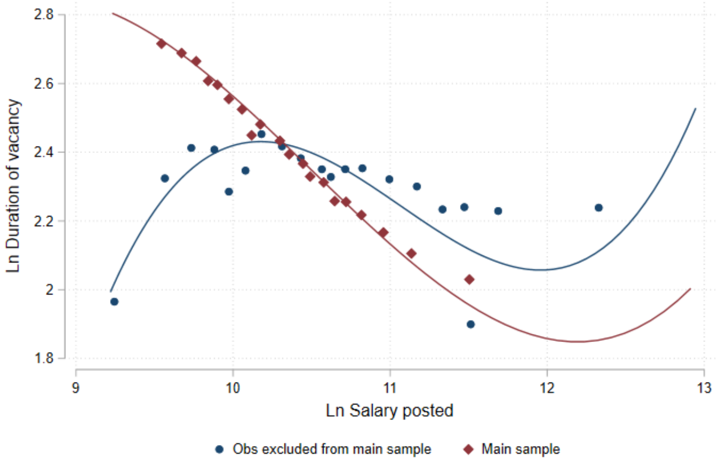

Notes. The binned scatterplots control for job and date-by-TTWA FE (following Cattaneo et al., 2022). The sample corresponds to column 3 in Table C1. "Main sample" excludes the 1% tails of residualizes wages, and "Obs excluded" denotes observations excluded by trimming (about 0.4 million adverts). The linear slope for the main sample is −0.4 and for the excluded sample is −0.07. The plot omits observations bunched at zero residualized wage for better visualization (slope estimates are similar when included).

Figure C7: Baseline estimates for alternative levels of wage trimming

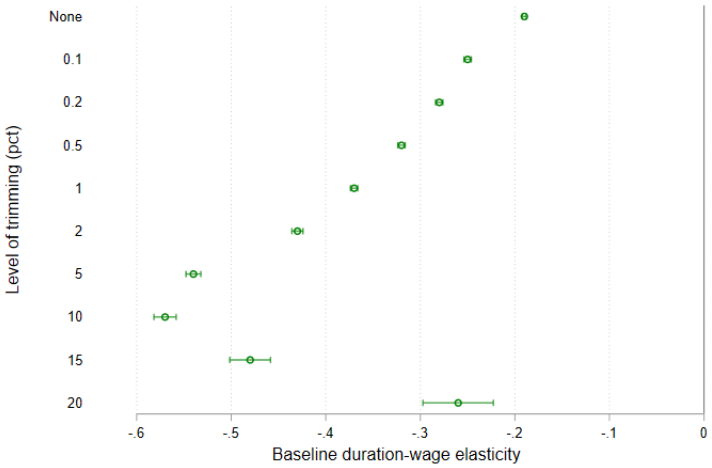

Notes. The figure shows coefficients from separate regressions of log vacancy duration on log wages, controlling for job and date-by-TTWA FE (see specification 2). The sample corresponds to column 3 of Table C1. Alternative levels of trimming refer to the tail of the distribution of wages residualized with respect to job and date-by-TTWA FE. Horizontal bars indicate 95% confidence intervals.

Figure C8: Alternative specifications

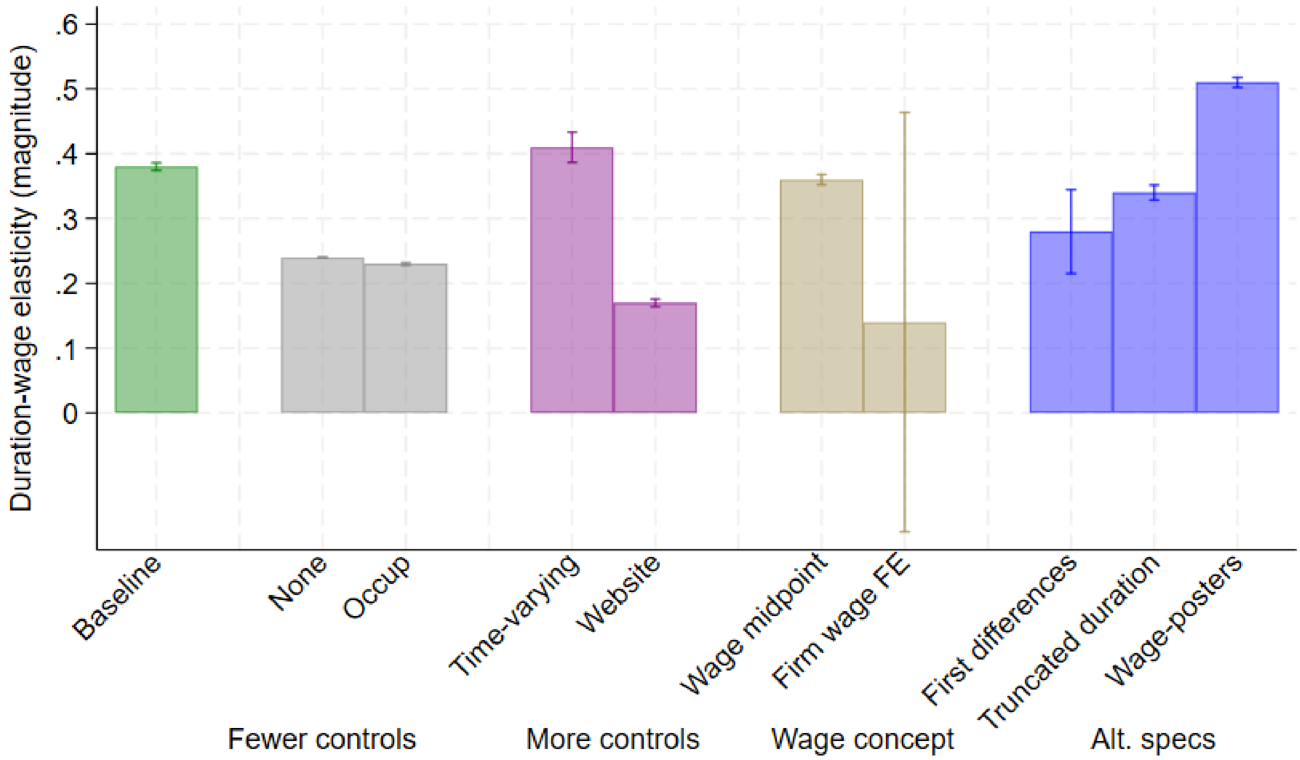

Notes. The figure shows coefficients from separate regressions of log vacancy duration on log wages (see specification 2). The sample corresponds to column 3 of Table C1. All specifications include date-by-TTWA FE. Baseline estimates replicate those in column 2 of Table C6. "Fewer controls" refer to no controls ("None") and 4-digit occupation controls ("Occup"). "More controls" refer to (in addition to job FE): firm-level employment growth, number of employees, sales and vacancies ("Time-varying"); FE for the original posting website ("Website"). "Wage concept" refers to: mid-point of the posted wage-range ("Wage midpoint"), as opposed to the maximum; firm wage FE, obtained from a regression of posted wages on firm FE, TTWA, 4-digit industry, 4-digit occupation, controls for benefits and wage concept ("Firm wage FE"). Alternative specifications ("Alt. specs") include a regression on first differences; a censored duration regression that truncates vacancy duration at 3 weeks ("Truncated duration"); and the baseline specification restricted to the sample of jobs where all vacancies have wages posted ("Wage-posters").

Figure C9: Distribution of wage-change events

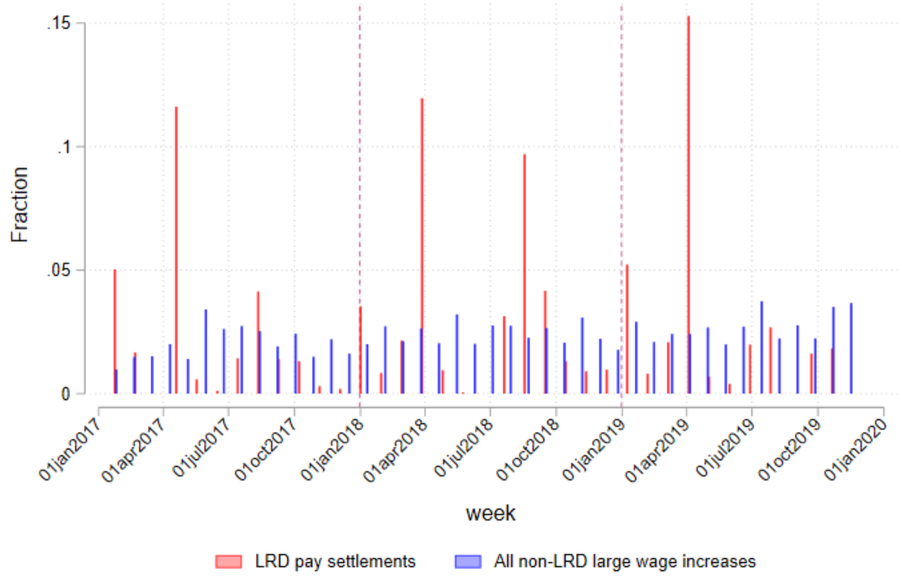

Notes. The figure shows the fraction of wage-change events taking place each week during 2017-2019. Red bars refer to pay settlements in the LRD database; blue bars refer to internally-defined wage-change events (see text for full definitions).

Figure C10: Magnitude of wage changes in the LRD database

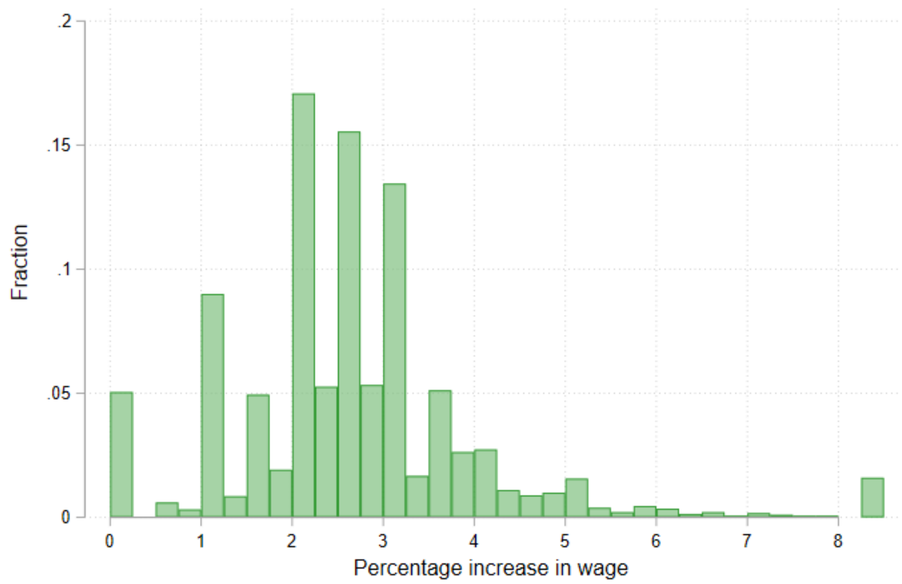

Notes. The figure shows the distribution of wage changes in the LRD pay-settlement database during 2017-2019. The right-most bar corresponds to wage increases above 8%.

Figure C11: Estimates based on external information on pay settlements: Alternative levels of wage trimming

(a) IV estimates

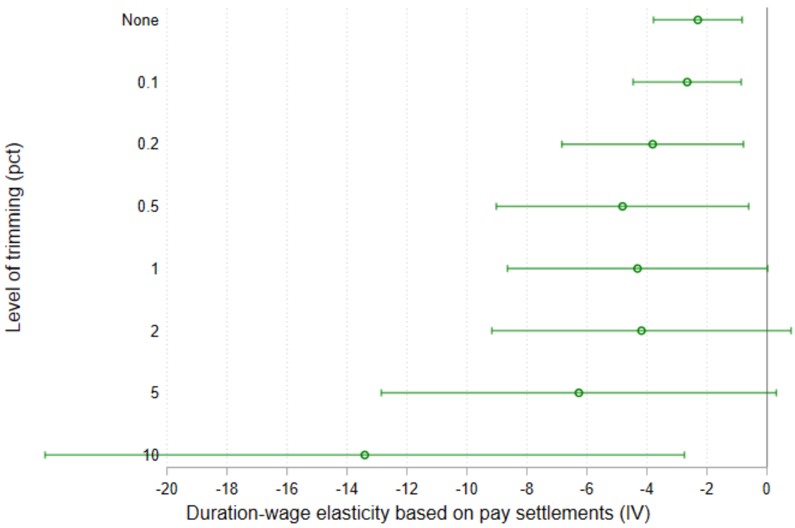

(b) First stage estimates

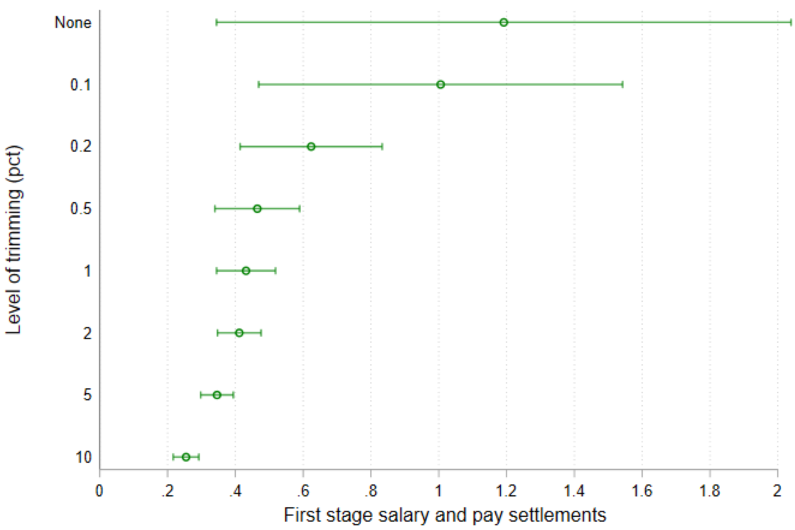

Notes. The figure shows coefficients from separate IV and first-stage regressions, using pay settlements in the LRD database as instruments for posted wages. The specification corresponds to column 3 in Table 2. Alternative levels of trimming refer to the tail of the distribution of wages residualized with respect to job and date-by-TTWA FE. Horizontal bars indicate 95% confidence intervals.

Figure C12: Wage changes before and after internally-defined wage events

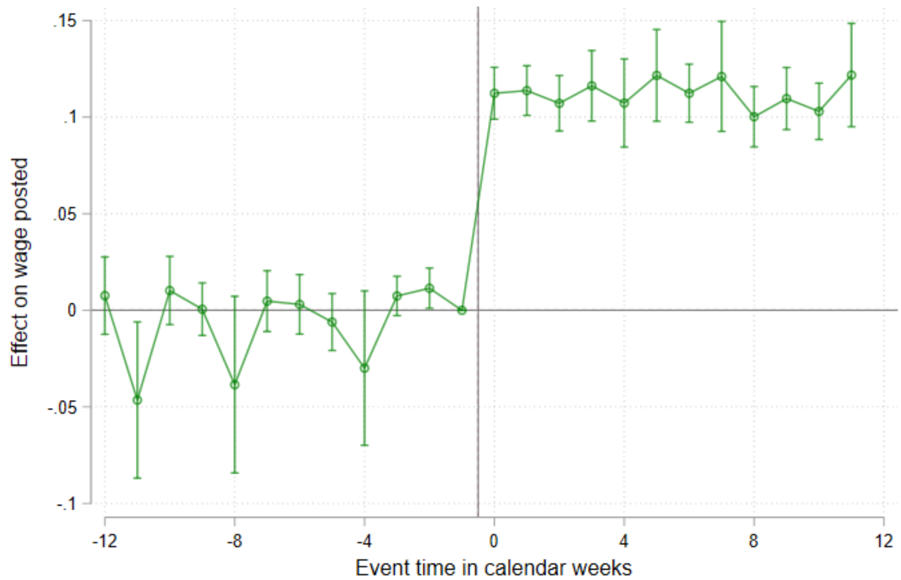

Notes. The figure shows coefficients from a regression of (log) wages on weekly event-time effects, controlling for job and date-by-TTWA FE. The sample includes cleaned Adzuna vacancy data for 2017–2019, restricted to firms experiencing an internally-defined wage-change event and corresponding control firms. 0 indicates the time of the event.

Figure C13: Event study estimates: Controlling for treatment-by-calendar time effects

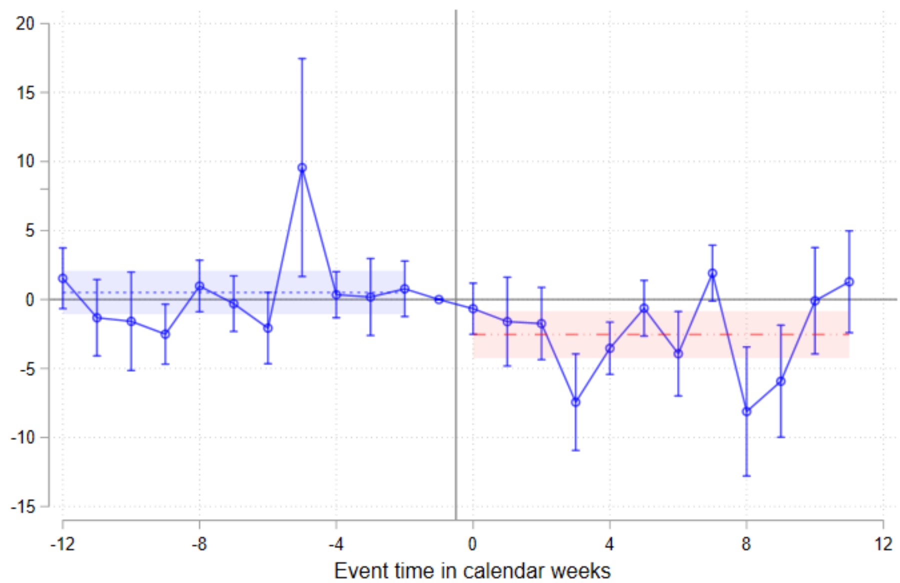

Notes. The figure shows robustness on the event-study results shown in Figure 1, by additionally controlling for event time-by-week FE, thus estimates only leverage the magnitude of the firm-level wage changes. All other details are as in Figure 1.

Figure C14: Event study estimates: Wage events defined on alternative criteria

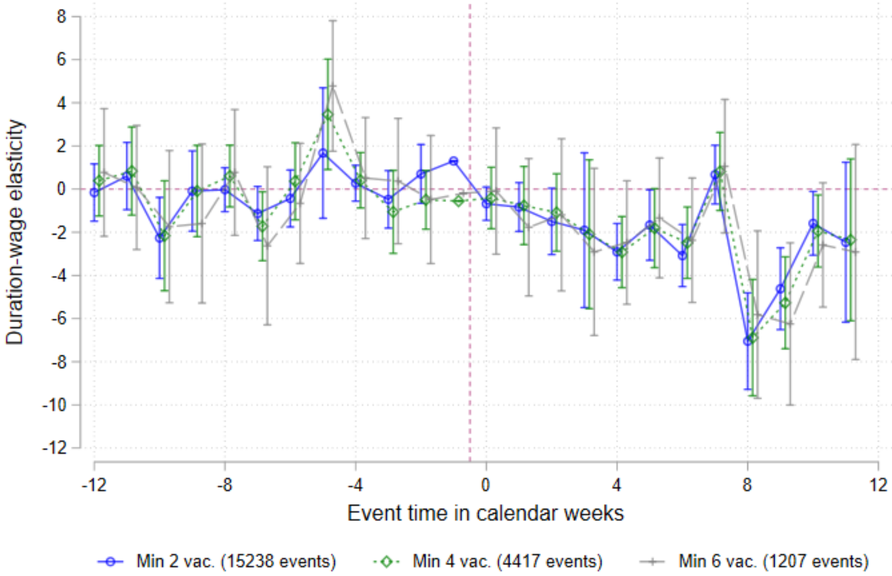

Notes. The figure shows robustness on the event-study coefficients shown in figure 1, by defining events under alternative criteria on the number of vacancies available in the event window. Imposing a minimum of 2 vacancies (1 pre and 1 post) yields 15,238 wage-change events in our sample, according to the definition of Section 5.1; our main specification, imposing a minimum of 4 vacancies, yields 4,417 events; imposing a minimum of 6 vacancies yields 1,207 events. All other details are as in Figure 1.

Figure C15: Event study estimates: Leave-one-out specification

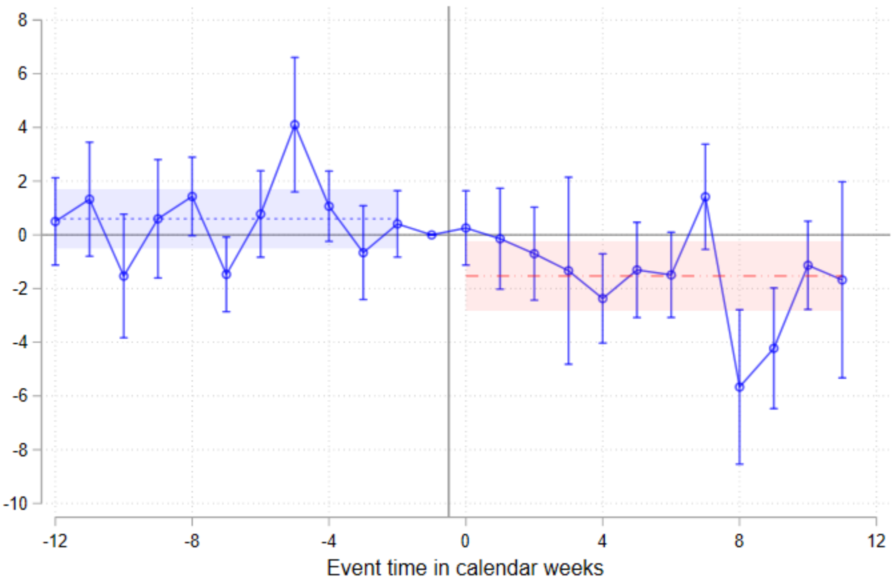

Notes. The figure shows robustness on the event-study results shown in Figure 1, by using as instrument the leave-one-out mean of firm-level wage changes. All other details are as in Figure 1.

Table C1: Descriptive statistics on estimation samples

Sample	All		Baseline	External		Internal	
	Raw	Clean		Treat	Control	Treat	Control
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Vacancies (th.)	50150	32524	21660	66	333	19	211
Jobs (th.)	20787	16538	6038	19	114	3.5	7.6
Wage (th., mean)	37.8	37.8	37.1	25.4	35.8	28.7	36.2
Wage (th., p50)	30	30	30	19.2	25	21	30
Duration (mean)	17.7	18	16.3	13.8	14.8	17.6	11.6
Duration (p50)	15	17	13	11	11	18	8.7
Occupation (%)							
High skill	52.5	53.6	52.7	31.4	47.4	27.1	47
Mid skill	23.8	23.9	24.6	19.9	25.2	24.3	27
Low skill	23.6	22.5	22.7	48.7	27.4	48.6	26

Notes. The Table describes alternative samples selected from the Adzuna database from 2017-2019. Column 1 refers to all adverts in the raw sample, excluding only those with missing values for company name, TTWA or title (about 2 million) or whose recorded duration does not match the number of observed vacancy stock (about 3 million). Column 2 includes adverts with non-missing wages. Column 3 includes the subsample with at least two adverts per job; this is used for baseline estimates in Table 1. Column 4 includes the subsample of firms that can be matched to the LRD pay-settlement database. Column 5 includes the corresponding control group (firms that are not matched to the LRD pay-settlement database and experience no large wage change over the sample period). The combined sample in columns 4 and 5 is used for estimates in Table 2. Column 6 refers to the sample of firms with internally-defined wage-change events (a change in firm-level wages of at least 5%, surrounded by 12 weeks on either side of nil wage changes). Column 7 refers to the corresponding control group (firms that do not experience a wage increase above 1% over the same 24-week interval). The combined sample in columns 6 and 7 is used for estimates in Table 3. The number of vacancies and jobs is measured in thousands; wages are measured in thousands GBP per year. Duration is measured in days. A job is defined by the combination of a job title, firm name, and TTWA. Occupations are grouped into: high-skill (managers, professionals and associates), mid-skill (administrative, and skilled trades and service occupations), and low-skill (operatives, sales and elementary occupations).

Table C2: Comparisons of estimates from: (1) OLS; (2) measurement error IV; (3) pay settlement IV

	(1)	(2)	(3)
First stage		0.073***	0.028***
		(0.006)	(0.006)
Reduced form		-0.186***	-0.163***
		(0.047)	(0.052)
Main equation	-0.655**	-2.562***	-5.727***
	(0.295)	(0.668)	(1.897)
F-stat		166.787	23.709
Treat sample	✓	✓	✓
Binary IV		✓	
Pay set. IV			✓
Vacancies	40498	40498	40498
Jobs	40498	40498	40498

Notes. The specification is in first differences, comparing a vacancy with the last vacancy posted for the same job. The outcome is the difference in vacancy duration, and the endogenous regressor is the difference in the wage posted. Column 1 reports OLS estimates, column 2 uses a binary indicator for a rise in the vacancy wage above 1% as an IV, and column 3 uses a binary indicator for a rise in the external wage settlement as an IV. "Treat sample" indicates that all regressions are restricted to the pay settlement sample (see column 4 of Table C1).

Table C3: Robustness on IV estimates: Additional time-varying labour market controls

	IV: External wage events						IV: Internal wage events	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
First stage	0.536***	0.016***	0.450***	0.014***	0.460***	0.015***	0.864***	0.481***
	(0.076)	(0.002)	(0.076)	(0.002)	(0.068)	(0.002)	(0.024)	(0.081)
Reduced form	-1.581	-0.076**	-2.091**	-0.062**	-1.761*	-0.078**	-1.090**	-1.765***
	(1.056)	(0.033)	(0.966)	(0.031)	(1.003)	(0.030)	(0.514)	(0.635)
Main equation	-2.951	-4.625**	-4.645**	-4.330*	-3.826*	-5.202**	-1.262**	-3.668***
	(2.045)	(2.092)	(2.330)	(2.220)	(2.226)	(2.061)	(0.594)	(1.224)
A-R CI	[-7.24,0.85]	[-9.01,-0.73]	[-9.90,-0.49]	[-9.16,-0.20]	[-8.49,0.67]	[-9.36,-1.20]	[-2.27,0.17]	[-6.33,-0.52]
F-stat	49.318	75.578	35.348	55.579	45.536	51.651	1320.398	35.490
Timing onl✓		✓		✓		✓		
Leave-one-out								Y
LM controls								
TTWA× week FE	✓	✓						
TTWA× month FE			✓	✓				
TTWA× occ. trends					✓	✓	✓	✓
TTWA× occ.× week FE							Y	Y
Vacancies	337239	337239	379084	379084	366237	366237	168291	148555
Jobs	112935	112935	126781	126781	122261	122261	10804	9061

Notes. See main text for specifications and samples corresponding to wage events from the external IV (Table 2) and the internal IV (Table 3). "LM controls" indicate labour market time-varying controls, where occupation is defined at the 1-digit level. A-R CI indicates the Anderson-Rubin confidence interval for IV estimates. Standard errors are reported in brackets.

Table C4: IV estimates: Alternative control samples

	IV: External wage events			IV: Internal wage events		
	(1)	(2)	(3)	(4)	(5)	(6)
First stage	0.660**	0.393***	0.414***	0.867***	0.778***	0.812***
	(0.265)	(0.062)	(0.062)	(0.048)	(0.116)	(0.077)
Reduced form	-2.737**	-1.700*	-1.960**	-0.912**	-2.172***	-2.057***
	(1.257)	(1.027)	(0.997)	(0.369)	(0.798)	(0.658)
Main equation	-4.145**	-4.329	-4.732*	-1.051**	-2.791***	-2.533***
	(1.696)	(2.703)	(2.520)	(0.434)	(1.063)	(0.808)
A-R CI	[.,-0.18]	[-10.00,0.92]	[-10.20,0.16]	[-1.89,-0.13]	[., .]	[.,0.06]
F-stat	6.196	39.792	44.541	320.973	44.685	112.017
Job FE	✓	✓	✓	✓	✓	✓
Location trends	✓	✓	✓			
Week × TTWA FE				✓	✓	✓
Control sample	None	P-score	N-N	None	P-score	N-N
Vacancies	65869	445677	128618	19066	1308437	40246
Jobs	19354	150107	49867	3955	248410	13036

Notes. See main text for specifications corresponding to wage events from the external IV (Table 2) and the internal IV (Table 3). Columns in this Table refer to alternative control samples. Samples in columns 1 and 4 ("None") only includes treated firms according to external and internal wage changes, respectively, i.e. no control firms are included, and estimates only leverage the magnitude and timing of the wage event as an IV. Samples in columns 2 and 5 ("P-score") include matched controls (using wages, benefits and location-specific trends as covariates) and estimates are based on propensity-score weights. Samples in columns 3 and 6 ("N-N") include matched controls obtained with nearest neighbor matching. The number of vacancies and jobs are reported as weighted counts. A-R CI indicates the Anderson-Rubin confidence interval for IV estimates, where a missing bound indicates an unbounded interval on that side. Standard errors are reported in brackets.

Table C5: Further robustness on IV estimates

	IV: External wage events			IV: Internal wage events			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
First stage	0.464***	0.547***		0.824***	0.369***	0.886***	0.693***
	(0.064)	(0.072)		(0.057)	(0.042)	(0.050)	(0.112)
Reduced form	-2.255**	-2.039**	-2.352**	-2.675***	-2.681***	-3.114***	-2.933***
	(0.931)	(0.946)	(1.108)	(0.793)	(0.953)	(0.930)	(0.871)
Main equation	-4.859***	-3.727**		-3.246***	-7.259***	-3.514***	-4.231***
	(2.149)	(1.816)		(0.999)	(2.678)	(1.122)	(1.411)
A-R CI	[-9.36,-0.86]	[-7.53,-0.35]		[-6.10,-0.68]	[-11.80,-0.58]	[-7.73,-0.63]	[.,-0.93]
F-stat	52.970	57.970		209.071	76.935	315.940	38.627
Alt. wage		✓			✓		
Controls							
Job/event FE	✓	✓	✓	✓	✓	✓	✓
Location trends	✓	✓	✓				
Week × TTWA FE				✓	✓	✓	✓
Vac. website FE	✓			✓			
Incl. no-wage ads			✓				
Min. 10 vacancies						✓	✓
Leave-one-out							✓
Vacancies	389167	365049	411204	227704	165891	224484	207325
Jobs	130297	121570	137919	11210	10153	10619	9795

Notes. See main text for specifications corresponding to wage events from the external IV (Table 2) and the internal IV (Table 3). Vacancy website FE refer to the original postiong website. "Alt wage" means that the middle of the posted wage range is used (as opposed to the top of the range in the main specification). "Min 10 vacancies" indicates that internal wage events are defined based on a minimum of 10 vacancies (as opposed to 3 in the baseline specification). A-R CI indicates the Anderson-Rubin confidence interval for IV estimates, where a missing bound indicates an unbounded interval on that side. Standard errors are reported in brackets.

Table C6: IV robustness: Restrictions on potential vacancy duplicates

	IV: External wage events				IV: Internal wage events			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
First stage	0.492***	0.016***	0.441***	0.015***	0.828***	0.591***	0.813***	0.559***
	(0.064)	(0.002)	(0.063)	(0.002)	(0.062)	(0.127)	(0.084)	(0.165)
Reduced form	-2.151**	-0.073***	-2.231**	-0.086***	-2.346***	-2.112***	-2.529***	-2.154***
	(0.957)	(0.028)	(0.968)	(0.028)	(0.644)	(0.611)	(0.742)	(0.661)
Main equation	-4.375**	-4.692**	-5.055**	-5.750***	-2.833***	-3.571***	-3.109***	-3.852***
	(2.050)	(1.874)	(2.351)	(1.965)	(0.801)	(1.136)	(0.988)	(1.469)
A-R CI	[-8.67,-0.56]	[-8.62,-1.20]	[-10.10,-0.68]	[-9.87,-2.00]	[-5.71,-0.90]	[.,-1.20]	[.,-0.80]	[.,-1.23]
F-stat	58.883	88.949	48.937	74.814	177.255	21.712	94.666	11.496
Timing only		✓		✓				
Leave-one-out						✓		✓
Duplicates restriction								
Across employers	✓	✓			✓	✓		
Over time			✓	✓			✓	✓
Vacancies	341822	341822	233226	233226	182142	165386	96875	88030
Jobs	120661	120661	130005	130005	10576	9680	10611	9555

Notes. See main text for specifications corresponding to wage events from the external IV (Table 2) and the internal IV (Table 3). The restriction "across employers" removes any additional vacancy sharing the same title and location within the same week. The restriction "over time" removes any additional vacancy posted within 8 days of the closing of the previous vacancy, for a specific job (employer by title by location).

References

Carrillo-Tudela, Carlos, Hermann Gartner, and Leo Kaas, "Recruitment policies, job filling rates and matching efficiency," *Journal of the European Economic Association*, 2023, pp. 2413-2459.

Cattaneo, Matias D, Richard K Crump, Max H Farrell, and Yingjie Feng, "On binscatter," Discussion Paper 1902 .09608, arXiv, 2022.

Manning, Alan, *Monopsony in motion: Imperfect competition in labor markets*, Princeton: Princeton University Press, 2003.