Department of Economics Discussion Paper Series

Vacancy Duration and Wages

Ihsaan Bassier, Alan Manning and Barbara Petrongolo

Number 1020
August, 2023

Department of Economics
Manor Road Building, Manor Road
Oxford, OX1 3UQ

ISSN 1471-0498

Vacancy Duration and Wages

Ihsaan Bassier, Alan Manning and Barbara Petrongolo*

August 3, 2023

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 identifies duration elasticities by leveraging firm-level wage policies that are plausibly exogenous to hiring difficulties on specific job vacancies, and control 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, Houghton Street, London WC2A 2AE, UK (email: i.bassier@lse.ac.uk); Manning: Centre for Economic Performance and Department of Economics, London School of Economics, Houghton Street, London WC2A 2AE, UK (e-mail: 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 (e-mail: 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

The elasticity of labor supply to the wage is central to the assessment of the degree of competition in labor markets. In a dynamic labor market model, higher wages ease worker recruitment and retention, and the elasticities of the 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 elasticity of separations to the wage (see Sokolova and Sorensen, 2021), but evidence on the recruitment margin is more limited. The underlying complication is that it is typically impossible to characterise the set of prospective employers and associated wages in a worker's opportunity set. Our paper focuses on the elasticity of vacancy durations 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. Estimates in this range are considerably larger than existing estimates for the vacancy elasticity, as well as OLS estimates in our data, even when controlling for very 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. We argue that a valid identification strategy needs to adequately address potential biases arising from unobserved job, worker and firms characteristics, the likely correlation of firm-level wages with those of competitor firms, and the possible endogeneity of wage adjustments to perceived hiring difficulties.

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 Carrillo-Tudela, Gartner, and Kaas (2020), 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, Marinescu, and Steinbaum, 2022; Banfi and Villena-Roldan, 2019; Belot, Kircher, and Muller, 2018), as vacancies that offer higher wages and attract more applicants are expected to fill faster, and on the elasticity of recruitment (Dal Bó, Finan, and Rossi, 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

The key feature of monopsonistic labor markets is that the labor supply to an individual firm is not infinitely elastic, hence the wage elasticity of firm-level employment is often used to measure employer market power. In a dynamic labor market model, steady-state firm-level 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 thus given by the difference between the recruitment and the 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, Gartner, and Kaas (2020) 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 scant, as datasets containing information on the wage offered and the time to fill a vacancy are rare. Seminal estimates in this field by Faberman and Menzio (2018) are based on a one-off survey of the recruitment activities of US employers in the early 1980s. To relate vacancy duration to wages, Mueller et al. (2023) combine information on vacancies posted on the Austrian employment service platform with administrative data on starting wages paid to new hires by the posting firms. We complement these approaches with an analysis on a very large sample of job adverts in the UK, containing information on both vacancy duration and posted wages.

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 (unlike other providers, e.g. Burning Glass, now Lightcast, which measure the inflow), and are used by the UK's Office for National Statistics as an indicator of UK economic activity (for detail, see Office for National Statistics, 2023). The total stock of vacancies at the monthly level in Adzuna is on average 93% of the vacancy stock resulting from the separate ONS vacancy survey of businesses, which makes us confident that our dataset covers the vast majority of job adverts in the UK.

The dataset contains information on the name of the posting firm, the location, job title and wage, as well as a free-format job description. Each vacancy has a unique identifier, which allows us to link observations across weeks and measure its duration, i.e. the number of weeks it remains posted. Within a job vacancy, all characteristics stay constant across weeks.

Our sample covers vacancies first posted during 2017-2019, as later data are affected by pandemic-related restrictions. For vacancies posted in late 2019 we use 2020 data to measure their completed duration.[1] We organise the data into a collection of vacancy spells, defined by the first and last week when a vacancy is posted.

There are about 55 million vacancies advertised in our sample period; we exclude 3 million vacancies for which information on the date first posted is inconsistent with the weeks during which a vacancy is observed and additionally drop 2 million vacancies with missing information on the posting firm, location or job title. We observe wages for about two thirds of vacancies, a higher incidence than in other vacancy datasets (e.g. 16% in Burning Glass

[1]Virtually all vacancy durations are shorter than 2 months, thus this procedure does not extend our coverage to the pandemic period. As an additional note, the Adzuna records do not contain information for December 2019.

data, see Hazell et al. 2022). 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; the distribution of workers across occupation skill terciles is also similar (see columns 1 and 2 in Table A1). Banfi and Villena-Roldan (2019) estimate a higher elasticity of applications to the wage when remuneration is posted, as opposed to implicitly conveyed in the job description, consistent with the idea that worker expectations about wages may be inaccurate when they are not posted. Thus, estimates of the duration-wage elasticity in our sample may be somewhat larger than in the universe of vacancies. Following these selection criteria, we are left with a sample of 32.5 million vacancies.

We match the Adzuna data with some external data sources. First, we match firm names in Adzuna with a database of firm pay settlements collected by the Labour Research Department's (LRD) .We obtain a (fuzzy) match for about half of the LRD agreements, but only 71,000 Adzuna vacancies. Secondly, we (fuzzy) match job titles with 4-digit occupations, obtaining a match for two thirds of observations, and (fuzzy) match firm names with the Orbis database, containing information on industry and other firm characteristics, obtaining a match for about half the sample.

3.2 Descriptive evidence

We define a job j as any vacancy with the same job title, posting firm, and location, defined at the travel to work area (TTWA) level.[2] 70% of vacancies for which industry information can be matched are posted by recruitment 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 recruitment agencies. Vacancies are posted relatively evenly across the months of the year, though are concentrated on the first day within each month (see Figure A1). For most jobs j, there are usually only a few vacancies advertised over the sample period: 65% of jobs have only one vacancy,

[2]There are 228 TTWAs in the UK, defined by the ONS to be commuting zones.

corresponding to 40% of vacancy observations in the sample (see Figure A2).

Figure A3 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 up for a month. This spike remains in the subsample of vacancies posted directly by employers, suggesting that it may not be explained by monthly posting fees charged by recruitment agencies. In a robustness check, we estimate a truncated regression which censors vacancy duration at 3 weeks, to exploit variation that is least affected by the 4-week bunching.

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 much of the empirical literature on separations elasticities, whenever quits may not be distinguished from layoffs (see Sokolova and Sorensen, 2021). 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. The elasticity of observed vacancy duration can be written as

$$\frac{\partial \ln d}{\partial \ln w} = \frac{\partial \ln d_f}{\partial \ln w} + \gamma \left(\frac{\partial \ln d_u}{\partial \ln w} - \frac{\partial \ln d_f}{\partial \ln w} \right), \tag{2}$$

where $\frac{\partial \ln d_f}{\partial \ln w}$ and $\frac{\partial \ln d_u}{\partial \ln w}$ are the elasticities of the duration to fill and withdraw a vacancy respectively, and γ is the share of vacancies that are withdrawn. Studies that have access to information on withdrawals suggest that the second term in (2) – measuring the bias in estimated elasticities in our data – is unlikely to be large. Van Ours and Ridder (1992) and Mueller et al. (2023) find that 4% and 14% of their vacancy sample is withdrawn, respectively. This proportion is larger in M. Andrews et al. (2008) at 34%, but their sample of vacancies for teenagers is highly selected and withdrawal is indirectly inferred. Their estimates for the elasticity of filling or withdrawing a vacancy with respect to the wage are similar (0.09 and 0.04 respectively), and the posted wage does not significantly affect the incidence of withdrawals. Thus, two out of the three studies suggest relatively small incidence of withdrawals and the third suggests only a small difference in the estimated elasticity.

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,[3] and the lag between a job offer and the start of an employment spell (Davis et al., 2014; Van Ours and Ridder, 1992).

Wages are posted for 33 out of 55 million vacancies in our sample. Where a wage range is posted, the top of the range is used; most salaries are reported on an annual basis, otherwise we convert them on an annual equivalent if posted hourly or daily; we also exclude a very small number of annual salaries that are implausibly low (below £10,000) or high (above £1bn). We cross-check salary information from the dedicated vacancy field and from the free-format job description, estimating a correlation of 0.9 between the two. 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 A4 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 are obtained on the full analysis sample (see column 3 in Table A1) by regressing the log of completed vacancy duration on the log of the posted wage and a set of controls described below. Duration of search tends to be longer for more skilled workers, because the returns to match quality may be higher for specialized skills (Faberman and Menzio 2018, Amior 2019), and/or high-quality workers are relatively scarce.

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

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, defined by the interaction of a job title, firm identifier and location. Second, vacancy duration is likely to respond not just to own wage but also to the wage offered by competitor firms. As own and competitors' wages are likely to be affected by correlated shocks, we control for "market" fixed-effects, defined by the interaction of week-by-location effects. 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:

$$\ln d_{j,t} = \beta \ln w_{j,t} + \gamma_j + \alpha_{l,t} + \nu_{j,t}, \tag{3}$$

where the outcome variable $\ln d_{j,t}$ is the log duration of a vacancy for job j at calendar time t. The main coefficient of interest is β, denoting the duration elasticity with respect to the posted wage $w_{j,t}$. Job fixed-effects and week-by-location (TTWA) fixed effects are denoted by γ_j and $\alpha_{l,t}$, respectively.

Estimates are presented in Table 1. Column 1 estimates equation (3) on the full sample of vacancies and obtains an elasticity estimate of -0.195. This is significantly negative, unlike some other estimates reported in the literature (Faberman and Menzio, 2018; Mueller et al., 2023) but it is still substantially smaller than existing estimates of other margins of labor supply elasticity, see among others Sokolova and Sorensen (2021) and Hirsch et al. (2022). Column 2 introduces trimming in the wage data to reduce the impact of noise and measurement error: we residualize (log) wages by job and date-by-location fixed effects and drop 1% of observations at the extremes of the residuals' distribution. The estimated elasticity falls to -0.369. Figure A5 shows non-parametrically the relationship between residualized wages and duration on the trimmed sample and, for comparison, on the observations excluded by trimming. The negative relationship between wages and duration is much stronger on the trimmed sample, possibly because wage observations at the extremes of the distribution embody more measurement error, and/or because the elasticity may be different at the extremes.[4]

Column 3 in Table 1 includes other controls: whether the wage is posted as hourly (21% of vacancies), daily (8%), annual (53%) or unstated, and whether non-wage benefits are mentioned in the advert (13% of vacancies). The inclusion of these controls makes little difference to the estimated elasticity. Figure A7 considers additional controls, including firm-level employment growth and firm characteristics interacted with calendar time, and additional specifications based on firm fixed-effects (from a log wage regression that additionally controls for location, 4-digit occupation and 4-digit industry[5]), first-differences, or censored-duration regressions. Estimates range between -0.5 and -0.1.

The final two columns of Table 1 highlight the importance of controlling for job characteristics. When controlling for 4-digit occupation as opposed to job fixed effects (column 4), or firm-by-occupations fixed-effects (column 5), the elasticity is considerably smaller than in column 3, which controls for job fixed effects. This pattern is in line with results shown by Marinescu and Wolthoff (2020), who estimate that the elasticity of job applications to wages switches from negative to positive when occupation controls are replaced with more detailed job title controls.

5 Wage-change events

5.1 Research design

Our baseline regressions control for job fixed effects, absorbing the role of permanent job characteristics that may be systematically related to duration and wages. However, there remain concerns about reverse causality, i.e. firms may decide to post higher wages on a

[4]The estimated elasticity tends to fall with the level of trimming below 10% and rises thereafter (see Figure A6). We view a 1% trimming as a reasonable baseline choice.

[5]This analysis is restricted to vacancies on which the employer is observed, i.e. excluding vacancies from recruitment agencies.

certain job when they expect it to be harder to fill, leading to an upward bias in the estimated elasticity of duration to wages.

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, Kramarz, and Margolis, 1999). In their analysis, the use of the AKM firm component, as opposed to a worker's entry wage, lowers the elasticity estimates 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 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, just above -0.2 (whether on trimmed or original wage data, see the third last row in Figure A7). However, there may be concerns about the role of firm or worker unobservables that may be systematically correlated to the firm fixed effect. For example, this wage measure may reflect compensating differentials for permanent firm-level amenities or worker sorting – if high-wage firms attract workers with systematically different search durations on the employers' side. More generally, estimates of the elasticity of vacancy duration based on the AKM firm fixed-effect may be biased whenever the AKM assumptions do not hold, for example regarding the correlation between the firm effect and the worker-firm match effect (see the estimates and discussion by Bassier, Dube, and Naidu, 2022).

Our proposed 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 information on the firm’s name, the dates when wages were adjusted, and the associated, company-wide wage change. Figures A8 and A9 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 A1). We include control vacancies in this regression sample (as in the procedure recommended by Borusyak, Jaravel, and Spiess, 2022), i.e. vacancies at firms that do not feature in the LRD pay-settlement database and have no large 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 any length of time 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% (i.e. $\Delta \ln w_{f,t} \in [0.05, 0.5]$), and a surrounding 24-week interval without wage increases exceeding 1% (i.e. $\Delta \ln w_{f,t+h} < .01$ for $h \in [-12, 11]$ and $h \neq 0$). To limit the influence of any given job advert on the definition of a wage event, we restrict this sample to events involving at least three adverts.[6] 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 1,788 unique firm-level wage increases, covering 18,856 vacancies, corresponding to 3,461 jobs across 282 firms (see column 6 in Table A1). The number of vacancies used to define the firm wage change $\Delta \ln w_{f,t}$ varies widely across events. In the distribution of the wage-event sample size, the 25th percentile has 81 vacancies, but 16% are based on fewer than 10 vacancies. We therefore 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 change obtained on all adverts at time t, excluding job j. The leave-one-out sample has 1,694 unique firm-level wage increases, covering 18,288 vacancies, corresponding to 3,419 jobs across 247 firms.As control firms, we include in the estimating sample firms with a full 24-week span without any wage increase exceeding 1% (i.e. $\Delta \ln w_{f,t+h} < .01$ 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 A8 shows that the pay settlements are concentrated on certain dates, such as 1 April or 1 January, while the internally-defined settlements are more evenly spread out across the year.

Relative to the baseline sample (column 3 in Table 1), the LRD-matched vacancies cover a higher proportion of low-skill occupations, with lower wages and shorter durations (column

[6]There is a trade-off between setting a higher size threshold to define a wage event and the reduction in the sample size. We show robustness around this threshold in Table A3.

4). 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 pre- and post-event effects in an event-study design. The results support the hypothesis of parallel trends. Second, in the robustness analysis 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 covariates.

5.2 Estimates based on external information on pay settlements

Table 2 presents estimates of duration elasticities 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, simply controlling for job and location-specific time trends.[7] We obtain an estimate of about −0.1 on the raw wage data (column 1), falling 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 very similar results.

Columns 3 and 4 show results from IV estimates that use the external wage agreement as an instrument for $w_{j,t}$ in equation 3. Column 3 shows a first stage estimate close to 0.5. As expected, this is below 1, reflecting that wage changes in a given job may happen throughout the year, while the company-wide pay settlements happen once a year, and that these may not be fully binding for each job in a firm. 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

[7] 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.

report both conventional standard errors[8] and the conservative but more robust Anderson-Rubin confidence intervals for the IV estimates, to cater for potentially weak instruments(I. Andrews, Stock, and Sun, 2019; Lee et al., 2022). On the well-defined identification strategy of column 3, the elasticity estimate is about one order of magnitude larger than on the baseline estimate of column 2. This highlights the importance of exploiting largely exogenous wage events to avoid biases in OLS estimates.

One potential concern is that, while the timing of the wage settlement is largely predetermined in this sample, the magnitude of the wage change may be endogeneous to the firm's current hiring experience. Column 4 thus exploits the timing of the wage increase, using a step-wise dummy as a wage instrument. The first-stage estimate implies that firms who sign a pay settlement on average raise their wages by 1.4%, and the reduced-form estimate implies that they see a reduction in vacancy duration by 8.7%. The corresponding IV estimate is about -6.

We consider a number of robustness tests in Table A2. First, we restrict the regression sample to observations that are matched to the LRD, which addresses concerns of differential trends between treated and control firms and contamination by treated firms missing from the LRD database. This purely exploits changes in vacancy durations before and after a wage-change event. While the F-statistics on the first-stage are very small, Anderson-Rubin confidence intervals for the IV estimates exclude zero (columns 2 and 3). Next, we include matched controls (using wages, benefits, and location-specific trends as covariates) and report estimates based on propensity-score weights and nearest neighbour matching (which provide a better balance of covariates, see Goldschmidt and Schmieder, 2017; Roth et al., 2023). In all cases the IV estimates are in the same ballpark as the main estimates of Table A2. Figure A10 shows robustness on the extent of trimming, reporting very similar IV estimates except for especially high levels of trimming, for which confidence intervals are much larger.

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

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 higher 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 estimate of the elasticity of about -3.3.

One reasonable concern is that the strong first-stage partly reflects the definition of a wage-change event, which includes the wage change on job j. This may in turn be endogenous to firms' hiring conditions on the same job. In column 4 we introduce a slightly modified instrument, based on the leave-one-out mean of firm-level wage changes: $\Delta \ln w_{f,t} = \frac{1}{n_f - 1} \sum_{k \neq j}^{n_f} (\Delta \ln w_{k,t})$. As expected, the first-stage estimate is lower, at about 0.6, as is the F-statistics. The resulting IV estimate is still negative and highly significant at -4.3, with negative bounds. The IV estimates obtained are very robust to alternative selections of events. For example, in Table A3 we select events that involve at least ten vacancies (as opposed to three) and obtain closely comparable elasticities of -3.5 and -4.2, using overall means or leave-one-out means of firm-level wage changes as instruments, respectively. For additional robustness, Table A4 shows estimates on a sample that excludes control firms (column 2) or uses matched firms as controls (columns 3 and 4).

Overall, we find that duration elasticities are considerably larger when using cleaner research designs. This pattern is consistent with findings from the literature on the wage elasticity of separations, which estimates elasticities in the range -0.5 to -1.5 on cross-sectional specifications, but larger elasticities in the range -3 to -5 on better research designs (Dube, Giuliano, and Leonard, 2019; Sokolova and Sorensen, 2021).

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:

$$\ln d_{j,t+\tau} = \sum_{u=-12,u\neq-1}^{u=11} \beta_u \Delta \ln w_{f,t} \times 1\{\tau = u\} + \gamma_j + \alpha_{l,t} + \psi_\tau + \nu_{j,t+\tau}, \tag{4}$$

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 4 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. Figure A11 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. In contrast to typical event studies, the main regressor $\Delta \ln w_{f,t}$ is a continuous treatment dosage, as discussed in Callaway, Goodman-Bacon, and Sant'Anna (2021) and Chaisemartin et al. (2022). The standard errors are clustered at the firm level.

Figure 1 plots estimates of the event dummies β_u, denoting 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), suggesting that these large firm-level wage increases are not systematically related to prior dynamics in vacancy durations. The average post-period coefficient is negative and significant, which 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.

In the Appendix we implement several robustness checks. Figure A12 shows the main event study coefficients when including week by treatment fixed effects, so that the variation in treatment is purely from the magnitude of the firm wage increase; the pattern is similar, with a slightly larger elasticity point estimate. As the sample used for our event-study analysis is highly unbalanced, Figure A13 shows robustness with respect to the minimum number of vacancies used in each event. The pattern of estimated elasticities is very similar when using at least 2 or 6 vacancies to define a wage-change event (where the main analysis of Figure 1 imposes at least 4 vacancies). Finally, Figure A14 shows that the event-study estimates are similar when using leave-one-out firm-level wage changes.

6 Heterogeneity in duration elasticities

This section investigates heterogeneity in duration elasticities using the internally-defined measure of wage changes.[9] Figure 2 shows heterogeneity along three dimensions. First, as described in section 3, most vacancies (70%) are posted by recruitment agencies and we find that duration elasticities are larger for vacancies posted by agencies (-4.2) than for those posted by direct employers (-1.6). This possibly suggests that workers may more easily compare wages on similar jobs on agencies' websites, making behavior more sensitive to wage differences. This is an important, previously undetected, aspect of heterogeneity, given the rising importance of recruitment agencies. Second, we show that the magnitude of the duration-wage elasticity is 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. Third, estimates by skill levelgroups do not show significant differences.

Figure A15 in the Appendix reports results across the same categories as in Figure 2, obtained on the baseline specification of column 2 of Table 1. We have argued these estimated duration elasticities would be biased, but the heterogeneity analysis may still be informative if underlying biases are similar across groups. As expected given the much larger sample

[9] The sample based on external wage changes has relatively fewer wage-change events and we would lack power to investigate heterogeneous responses along various dimensions of interest.

size, all elasticities are precisely estimated. The wage elasticity is significantly larger for vacancies advertised by recruitment agencies and in tighter labor markets. They are also larger for higher skill occupations, though differences are very small.

7 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 ballbark as well-identified estimates of the separations elasticity from existing studies.

Table 1: **Baseline estimates of the duration-wage elasticity**

	(1)	(2)	(3)	(4)	(5)
Log wage	-0.195***	-0.369***	-0.383***	-0.215***	-0.188***
	(0.001)	(0.002)	(0.002)	(0.001)	(0.001)
Trimmed		Y	Y	Y	Y
Controls					
Date × TTWA FE	Y	Y	Y	Y	Y
Job FE	Y	Y	Y		
Additional controls			Y		
Occupation FE				Y	
Firm × Occupation FE					Y
No. vacancies (M)	21.62	21.24	21.24	14.48	14.47
No. jobs (M)	5.98	5.94	5.94	3.91	3.89

Notes. The table shows results from regressions of log vacancy duration on log wages and the indicated controls (see equation 3). The sample includes cleaned Adzuna 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 2-5 exclude observations in the 1% tails of residualized wages. Additional controls in column 3 include dummies for the wage concept (hourly, daily, annual) and for mentions of non-wage benefits in the advert. Occupation fixed-effects in columns 4 and 5 are at the 4-digit level. Standard errors are reported in brackets. The numbers of observations and vacancies are given in millions.

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

	(1)	(2)	(3)	(4)
First stage			0.465^{***}	0.014^{***}
			(0.064)	(0.002)
Reduced form			-2.238^{**}	-0.087^{***}
			(0.932)	(0.027)
Main equation	-0.110^{***}	-0.425^{***}	-4.816^{**}	-6.036^{***}
	(0.020)	(0.062)	(2.146)	(2.047)
A-R CI			[-9.32,-0.82]	[-10.30,-2.22]
F-stat			53.149	62.976
Job FE	Y	Y	Y	Y
Location trends	Y	Y	Y	Y
Trimmed		Y	Y	Y
Pay set. IV			Y	Y
No magnitude				Y
Vacancies	392773	389167	389167	389167
Jobs	130972	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 Adzuna 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 in calendar week interacted with location fixed effects. Columns 1 and 2 show baseline OLS estimates on this sample. Column 3 instruments the wage posted on a job by the firm-wide pay settlement from the LRD database. In column 4, the instrument is a dummy variable for a pay settlement. Trimming excludes 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)
First stage			0.824***	0.568***
			(0.057)	(0.125)
Reduced form			-2.675***	-2.421***
			(0.793)	(0.750)
Main equation	-0.074***	-0.124***	-3.247***	-4.261***
	(0.004)	(0.005)	(1.000)	(1.424)
A-R CI			[-6.13,-0.67]	[.,-1.04]
F-stat			209.010	20.680
Date X TTWA FE	Y	Y	Y	Y
Job FE	Y	Y		
Event FE			Y	Y
Trimmed		Y	Y	Y
Firm wage IV			Y	Y
Leave-one-out				Y
Vacancies	236592	232226	227704	210032
Jobs	12179	12167	11187	10257

Notes. The table shows coefficients from regressions of log vacancy duration on log wage, with fixed effects for job and date by travel-to-work area (TTWA). The sample includes cleaned Adzuna vacancy data for 2017–2019, restricted to firms experiencing an internally-defined wage-change event and corresponding control firms. Columns 1 and 2 show the baseline OLS specification on this sample. Columns 3 and 4 instrument the wage on the job advert by the mean and leave-one-out mean firm-level wage change, respectively. Event FE refer to weeks in each 24-week window around each event. Trimming excludes 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**

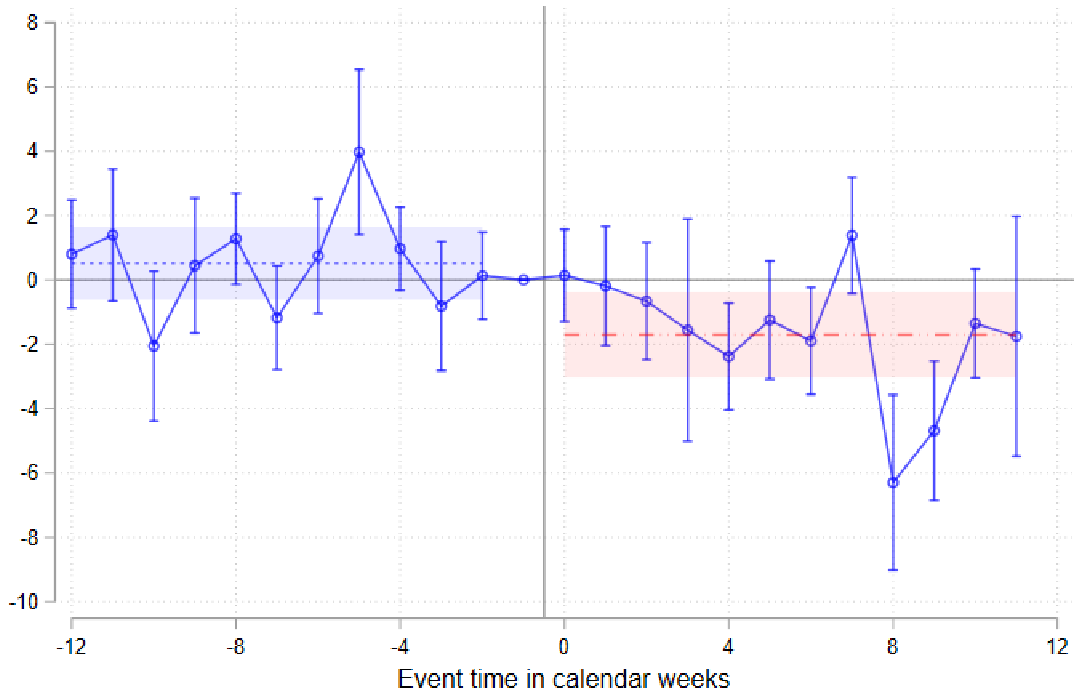

Notes. The figure shows weekly coefficients from an event-study regression of log vacancy duration on log firm-level wage changes, including fixed effects for jobs, event-time, and date by travel-to-work area (see equation 4). The sample includes cleaned Adzuna vacancy data for 2017–2019, restricted to firms experiencing an internally-defined wage-change event and corresponding control firms. 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. The (residualized) wage-change distribution is trimmed to exclude the 1% tails (excluding zero changes). Vertical bars and shaded areas represent 95% confidence intervals.

Figure 2: **Heterogeneity analysis on duration elasticities**

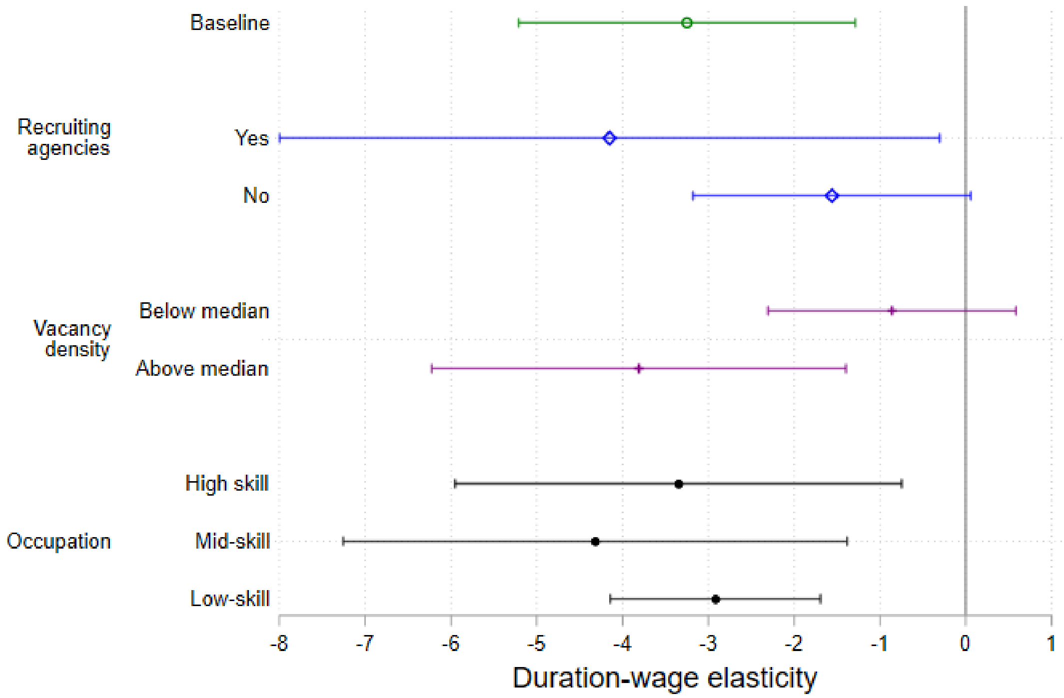

Notes. The figure shows coefficients from regressions of log vacancy duration on log wage, instrumented by internally-defined wage-change events, controlling for job and date by travel-to-work area fixed effects. The sample includes cleaned Adzuna vacancy data for 2017–2019, restricted to firms experiencing an internally-defined wage-change event and corresponding control firms (see column 3 of table 3). Each coefficient estimate is from a separate regression for the subsample indicated in the row header. The baseline sample has 11,187 jobs (matching column 3 of table 3). Recruiting agency status is determined by industry code merged from ORBIS data: recruits and non-recruits samples have 3,776 and 1,614 jobs respectively. Vacancy density is measured as the number of posted vacancies over regional employment, and samples have 998 and 8,792 jobs below and above median respectively. Occupations are drawn from 1-digit SOC2020 codes, and grouped in threes, i.e. 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 (1999). "High Wage Workers and High Wage Firms". In: *Econometrica* 67, pp. 251–333.

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

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

Andrews, Martyn, Steve Bradley, Dave Stott, and Richard Upward (2008). "Successful employer search? An empirical analysis of vacancy duration using micro data". In: *Economica* 75.299, pp. 455–480.

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

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

Banfi, Stefano and Benjamin Villena-Roldan (2019). "Do high-wage jobs attract more applicants? Directed search evidence from the online labor market". In: *Journal of Labor Economics* 37.3, pp. 715–746.

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

Belot, Michele, Philipp Kircher, and Paul Muller (2018). *How wage announcements affect job search-a field experiment*. Discussion Paper 3278661. SSRN.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess (2022). *Revisiting Event Study Designs: Robust and Efficient Estimation*. Discussion Paper 2826228. SSRN.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro HC Sant'Anna (2021). *Difference-in-differences with a continuous treatment*. Discussion Paper 2107.02637. arXiv.

Carrillo-Tudela, Carlos, Hermann Gartner, and Leo Kaas (2020). *Recruitment policies, job-filling rates and matching efficiency*. Discussion Paper 14727. Centre for Economic Policy Research.

Cattaneo, Matias D, Richard K Crump, Max H Farrell, and Yingjie Feng (2022). *On binscatter*. Discussion Paper 1902.09608. arXiv.

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

Dal Bó, Ernesto, Frederico Finan, and Martín A Rossi (2013). "Strengthening state capabilities: The role of financial incentives in the call to public service". In: *The Quarterly Journal of Economics* 128.3, pp. 1169–1218.

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

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

Dube, Arindrajit, Laura Giuliano, and Jonathan Leonard (2019). "Fairness and frictions: The impact of unequal raises on quit behavior". In: *American Economic Review* 109.2, pp. 620–63.

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

Falch, Torberg (2017). "Wages and recruitment: Evidence from external wage changes". In: *ILR Review* 70.2, pp. 483–518.

Goldschmidt, Deborah and Johannes F Schmieder (2017). "The rise of domestic outsourcing and the evolution of the German wage structure". In: *The Quarterly Journal of Economics* 132.3, pp. 1165–1217.

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

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

Lee, David S, Justin McCrary, Marcelo J Moreira, and Jack Porter (2022). "Valid t-ratio Inference for IV". In: *American Economic Review* 112.10, pp. 3260–90.

Marinescu, Ioana and Ronald Wolthoff (2020). "Opening the black box of the matching function: The power of words". In: *Journal of Labor Economics* 38.2, pp. 535–568.

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

Office for National Statistics (2023). *Using Adzuna data to derive an indicator of weekly vacancies: Experimental Statistics*. `https://tinyurl.com/pft5hkzm`. Accessed: 2023-07-07.

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

Sokolova, Anna and Todd Sorensen (2021). "Monopsony in labor markets: A meta-analysis". In: *ILR Review* 74.1, pp. 27–55.

Van Ours, Jan and Geert Ridder (1992). "Vacancies and the recruitment of new employees". In: *Journal of Labor Economics* 10.2, pp. 138–155.

A Additional Figures and Tables

Figure A1: **Distribution of vacancy postings over time**

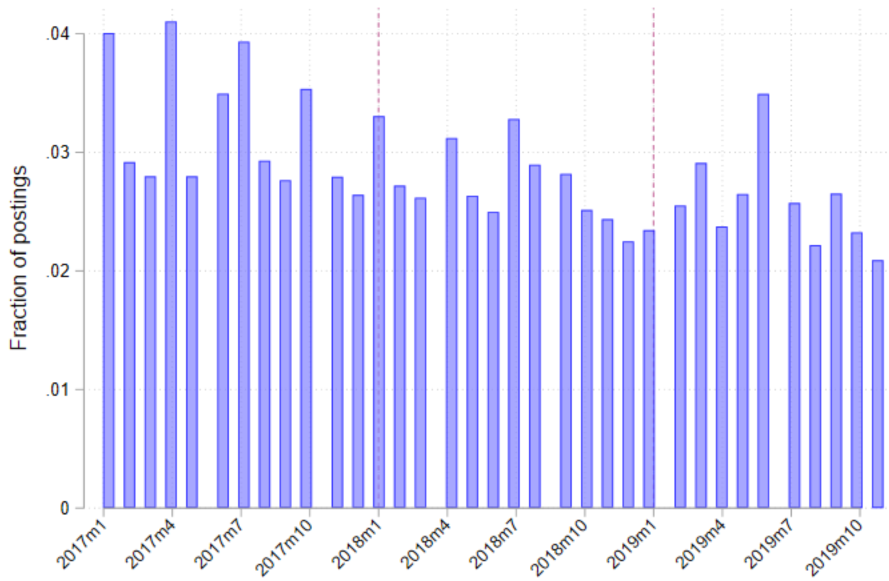

Notes. The graph plots the fraction of vacancy postings in each week during the sample period. The sample includes cleaned Adzuna vacancy data for 2017–2019, restricted to advert with non-missing wages (corresponding to column 2 in Table A1).

Figure A2: **Distribution of the number of vacancies per job**

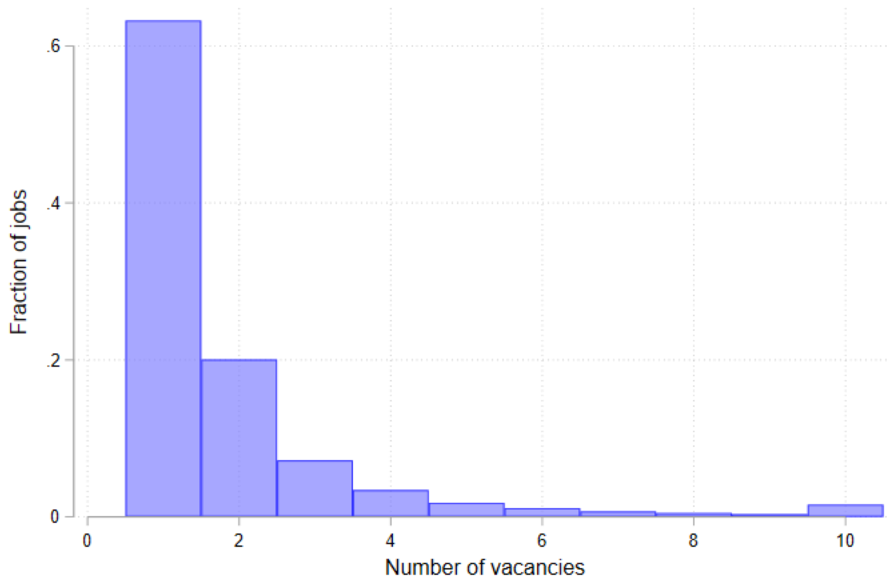

Notes. The graph plots the fraction of jobs with a certain number of adverts posted over the sample period. The sample includes cleaned Adzuna vacancy data for 2017–2019, restricted to advert with non-missing wages (corresponding to column 2 in Table A1). A job is defined by the combination of a job title, firm name, and location (TTWA). The final bar refers to jobs with 10 or more job adverts.

Figure A3: **Distribution of vacancy duration**

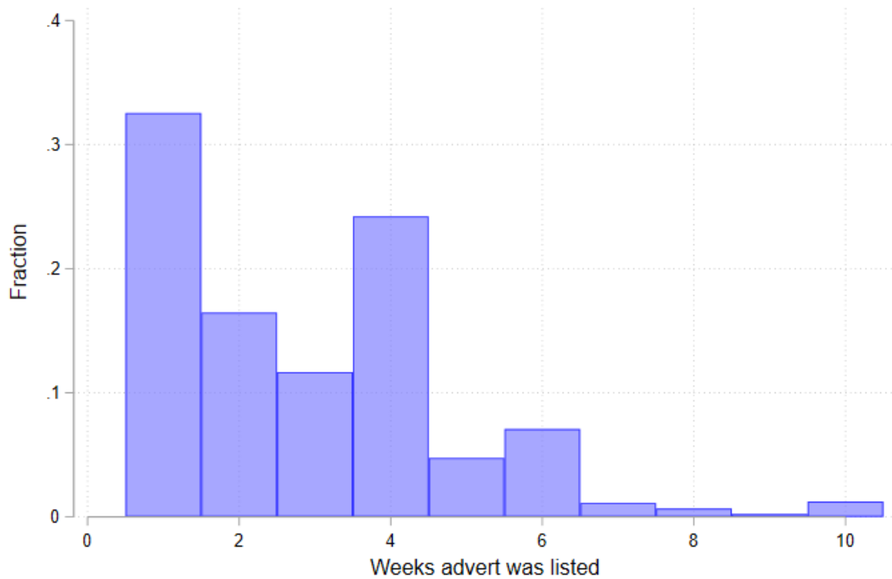

Notes. The graph plots the fraction of vacancies by the number of weeks they are posted. The sample includes cleaned Adzuna vacancy data for 2017–2019, restricted to advert with non-missing wages (corresponding to column 2 in Table A1). The final bar refers to durations of 10 weeks or longer.

Figure A4: **Wages in Adzuna and ASHE data**

(a) Levels

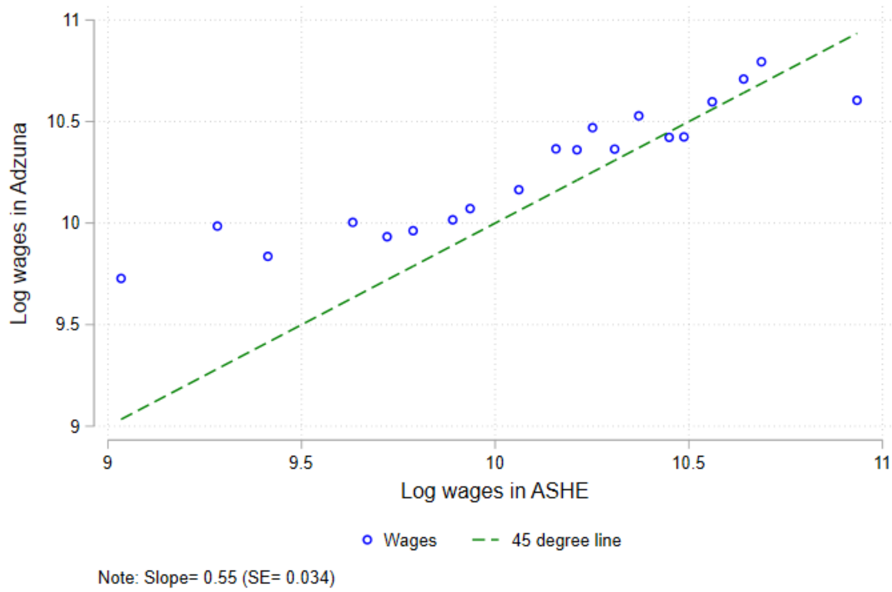

(b) Differences

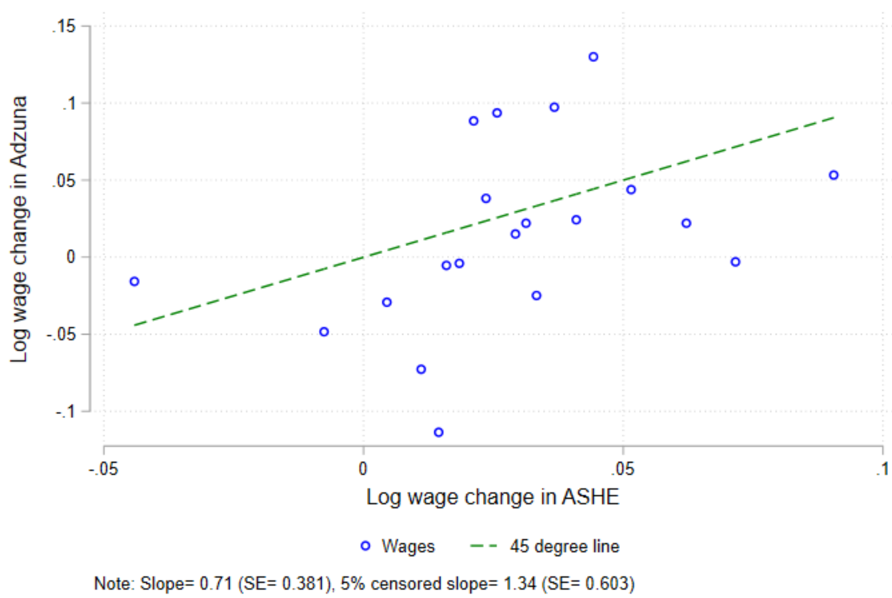

Notes. The figure shows a binned scatter plot of occupation-level wages (3-digit) in Adzuna and ASHE data. Panel (a) compares log wages across the two datasets and panel (b) plots the within-occupation annual wage change. The Adzuna sample includes cleaned vacancy data for 2017–2019, restricted to advert with non-missing wages (corresponding to column 2 in Table A1). The ASHE is an employer-based survey, covering a 1\% random sample of employee jobs in the UK. Data are weighted by the number of workers in each cell.

Figure A5: **Vacancy durations and wages**

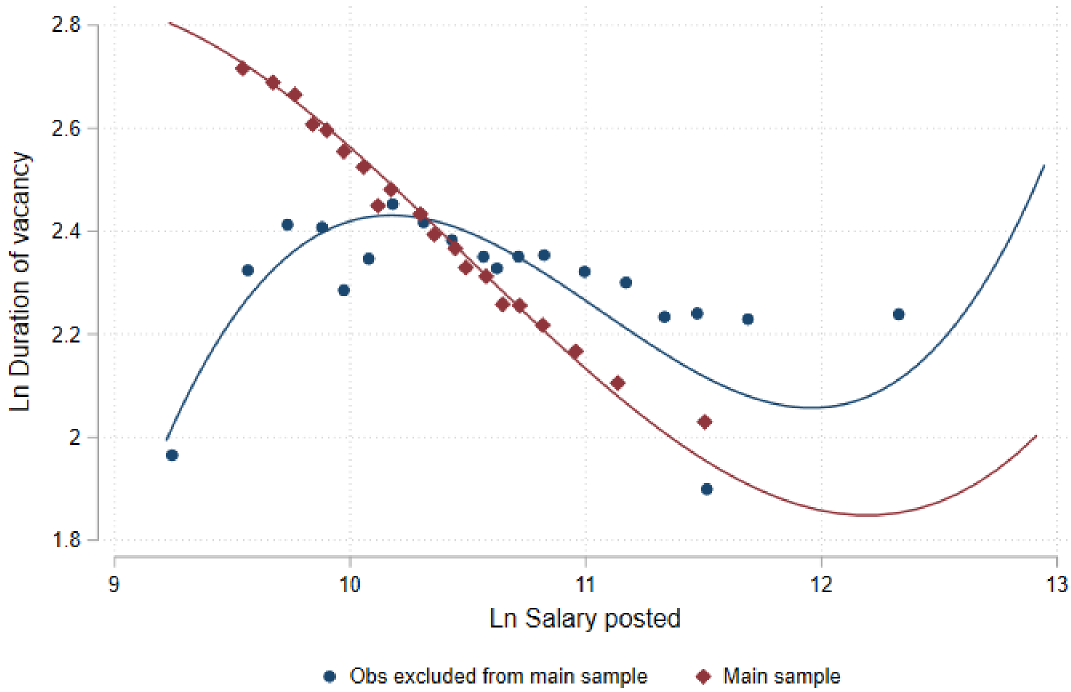

Notes. The figure shows a binned scatterplot of (log) vacancy by (log) posted wages, controlling for job and date-by-location fixed effects (adjusting for controls following Cattaneo et al., 2022). The sample includes cleaned Adzuna vacancy data for 2017–2019, restricted to jobs with at least two adverts (corresponding to column 3 in Table A1). The two scatter plots refer, respectively, to the main sample (where trimming excludes the 1% tails of residualizes wages) and the sample of observations excluded with 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 (the slope estimates is similar when included).

Figure A6: **Baseline estimates for alternative levels of wage trimming**

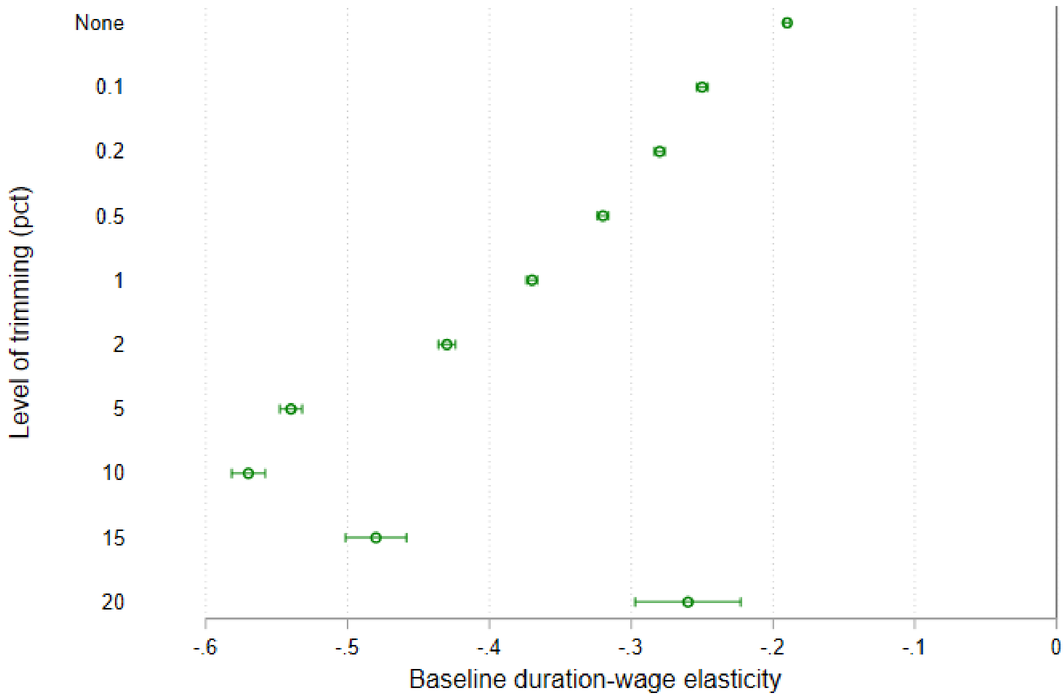

Notes. The figure shows coefficients from separate regressions of log vacancy duration on log wages, controlling for job and date-by-location fixed effects (see specification 3). The sample includes cleaned Adzuna vacancy data for 2017–2019, restricted to jobs with at least two adverts (corresponding to column 3 in Table A1). Estimates refer to alternative levels of wage trimming: percentages indicate the extent of trimming on each tail of the distribution of wages residualized with respect to job and date-by-location fixed effects.

Figure A7: **Alternative specifications**

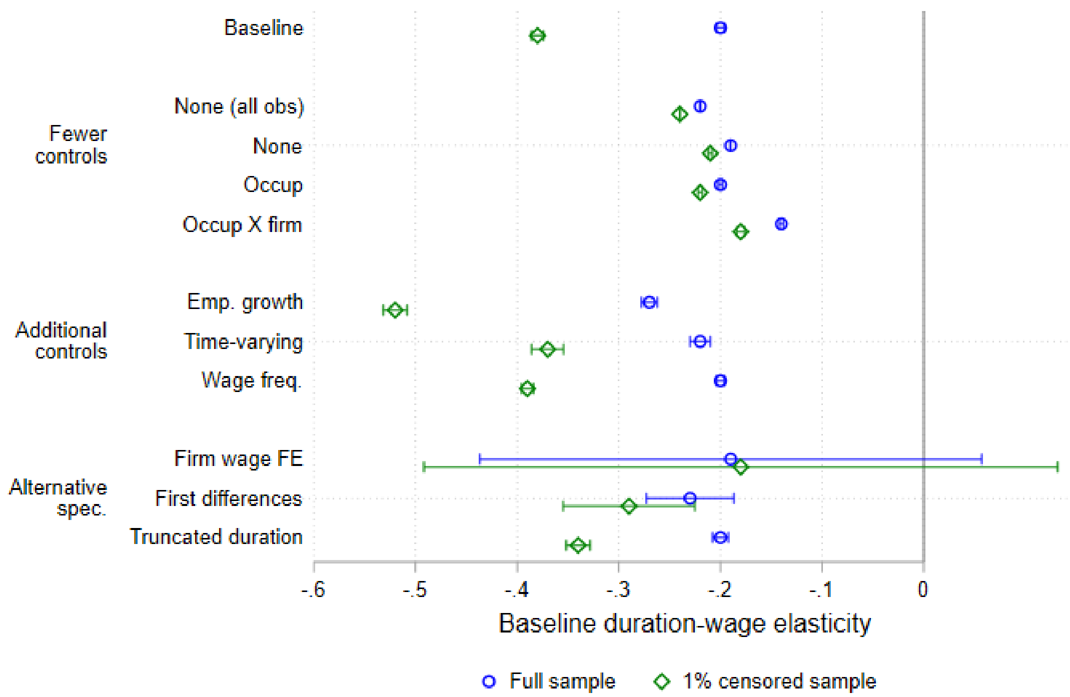

Notes. Each estimates is from a separate regression of (log) duration on (log) wages (see specification 3). The sample includes cleaned Adzuna vacancy data for 2017–2019, restricted to jobs with at least two adverts (corresponding to column 3 in Table A1), except row 2 ("None (all obs)"), which refers to the full sample (column 2 in Table A1). The baseline estimate coincides with that of column 2 in Table A5. All specifications with fewer controls include date-by-location FE. Additional controls refer to: firm-level employment growth; time-varying firm characteristics (number of employees, sales and vacancies); and indicators for wage concept (annual, weekly, hourly). Alternative specifications use: the firm FE as a regressor – obtained from a regression of posted wages on firm fixed-effects, location, 4-digit industry, 4-digit occupation, controls for benefits and wage concept (Firm wage FE); a specification in first differences; and a censored duration regression that truncates vacancy duration at 3 weeks.

Figure A8: **Distribution of wage-change events**

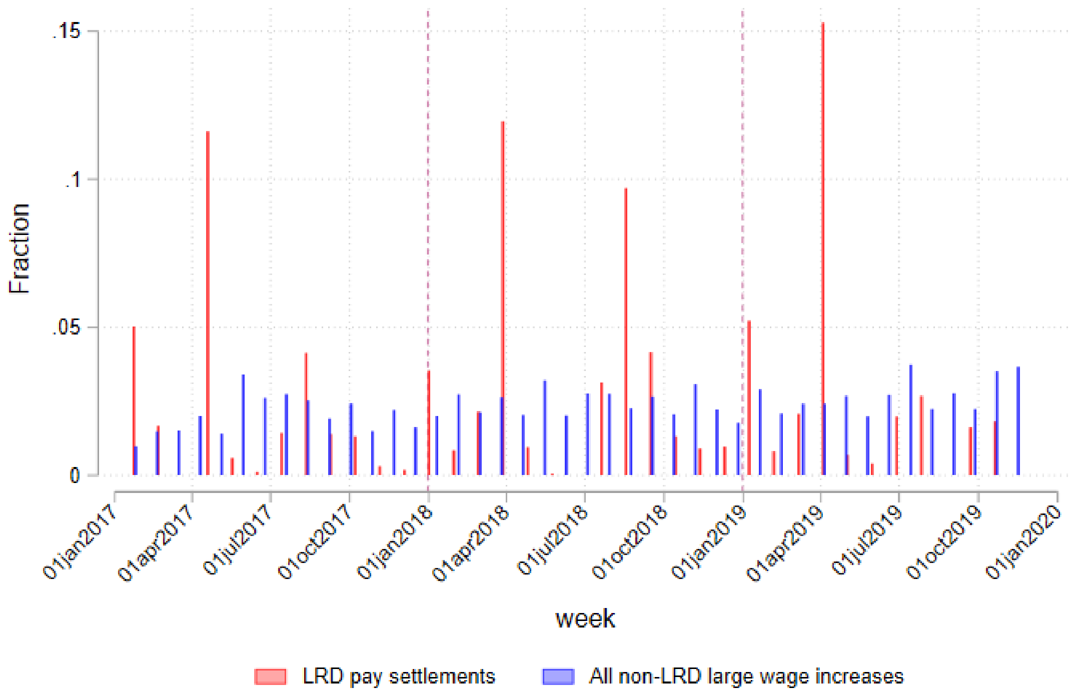

Notes. The graph shows that fraction of wage-change events taking place each week in the sample period. Red bars denote pay settlements in the LRD database; blue bars denote internally-defined wage-change events.

Figure A15: **Heterogeneity in duration-wage elasticity**

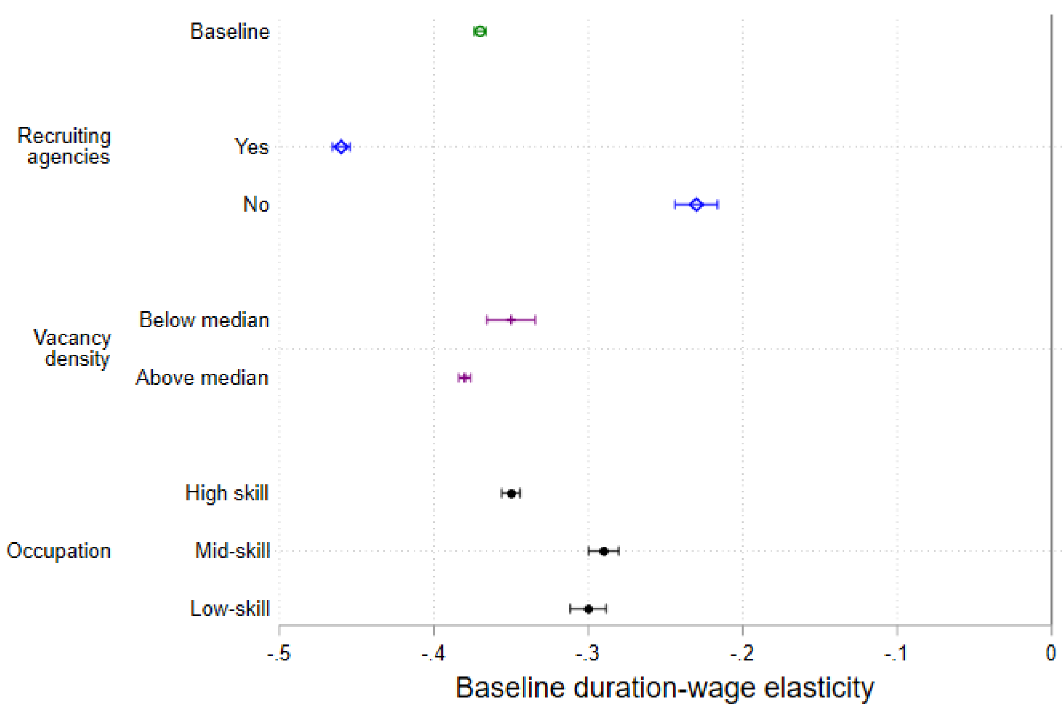

Notes. The figure shows coefficients from regressions of log vacancy duration on log wage, controlling for job and date-by-location fixed effects. The sample includes cleaned Adzuna vacancy data for 2017–2019, restricted to jobs with at least two vacancies (see column 2 of Table 1). Each coefficient

Figure A9: **Magnitude of wage changes in the LRD database**

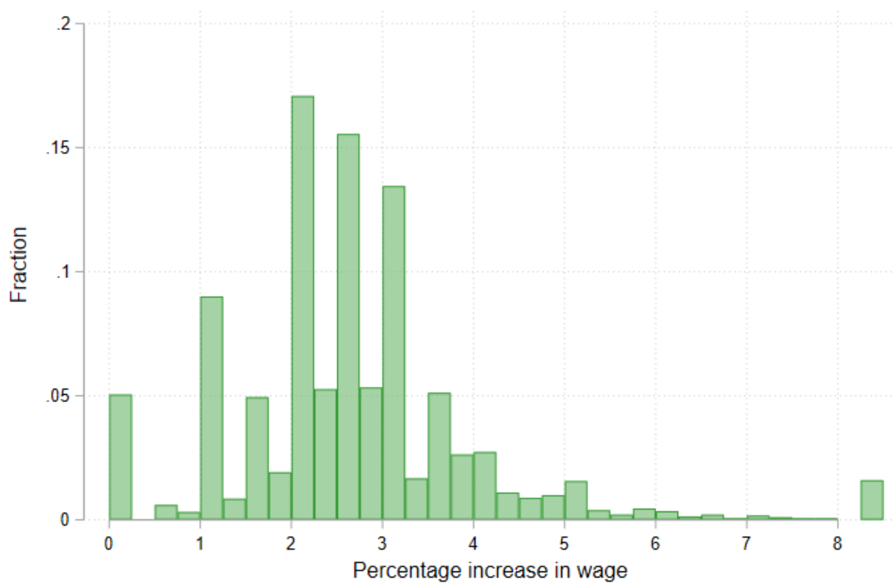

Notes. The graph plots the distribution of wage changes in the LRD pay-settlement database during 2017 to 2019. The final bar corresponds to wage increases above 8%.

Figure A10: **Estimates based on external information on pay settlements: Alternative levels of wage trimming**

(a) IV estimates

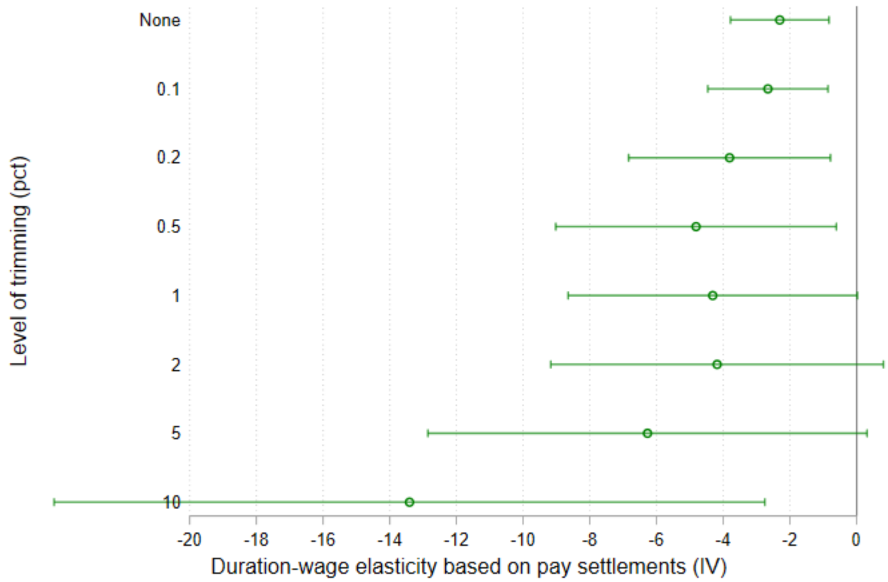

(b) First stage estimates

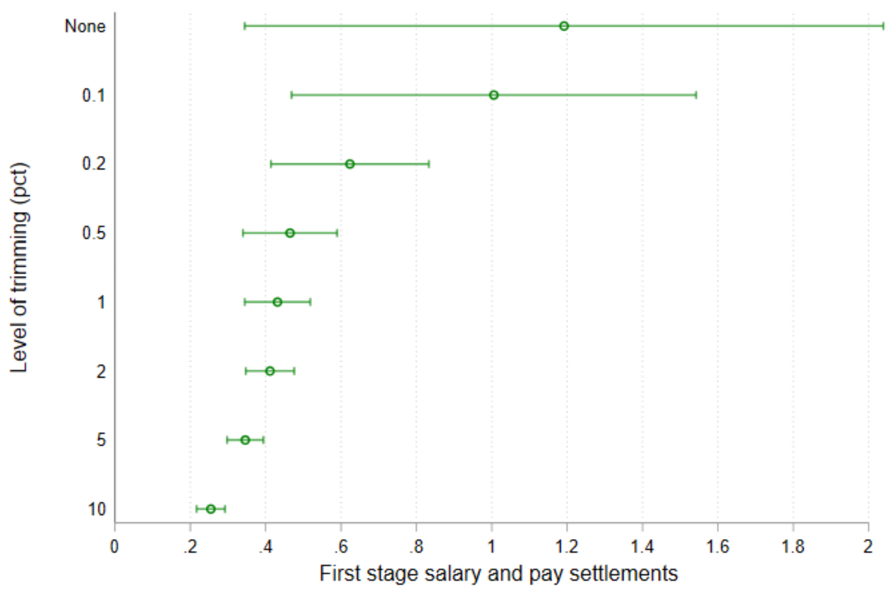

Notes. The figure shows coefficients from separate IV and first-stage regressions, using pay settlements in the LRD database as instruments for posted wages. The specification corresponds to column 3 in Table 2. The figure shows robustness of the IV estimates of log vacancy duration on log wage shown in Table 2. The sample includes cleaned Adzuna vacancy data for 2017–2019, restricted to firms that can be matched to wage agreements in the LRD database and corresponding control firms. Estimates refer to alternative levels of wage trimming: percentages indicate the extent of trimming on each tail of the distribution of wages residualized with respect to job and date-by-location fixed effects. Bars indicate 95% confidence intervals.

Figure A11: **Wage changes before and after internally-defined wage events**

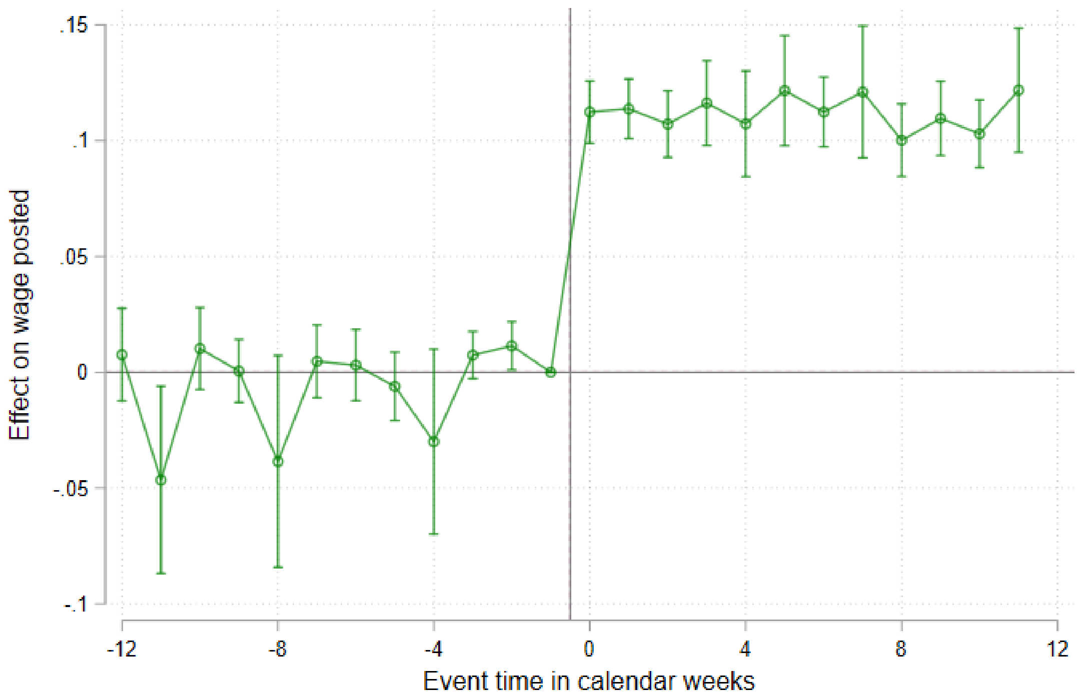

Notes. The figure shows coefficients from a regression of (log) vacancy wages on weekly event-time effects, controlling for job and date-by-location fixed effects. The sample includes cleaned Adzuna vacancy data for 2017–2019, restricted to firms experiencing an internally-defined wage-change event and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

Figure A12: **Event study estimates: Controlling for treatment-by-calendar time effects**

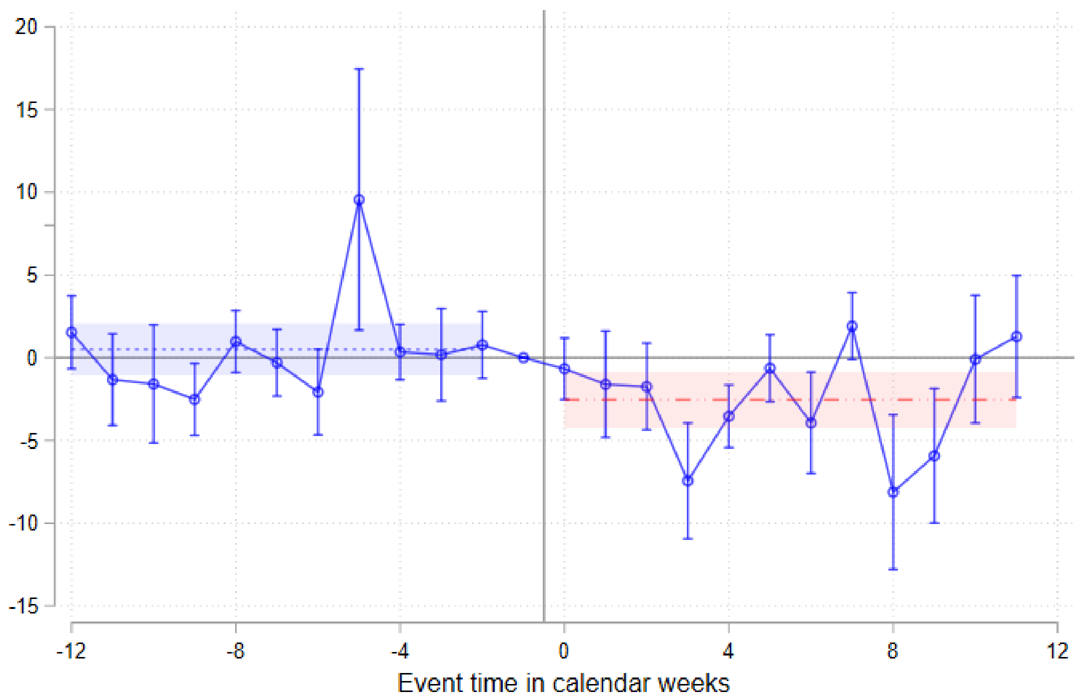

Notes. The figure shows robustness on the event-study results shown in Figure 1, by additionally controlling for event time-by-week fixed effects, so that coefficients are based purely on the magnitude of the firm wage policy changes. All other details are as in Figure 1.

Figure A13: **Event study estimates: Wage events defined on alternative criteria**

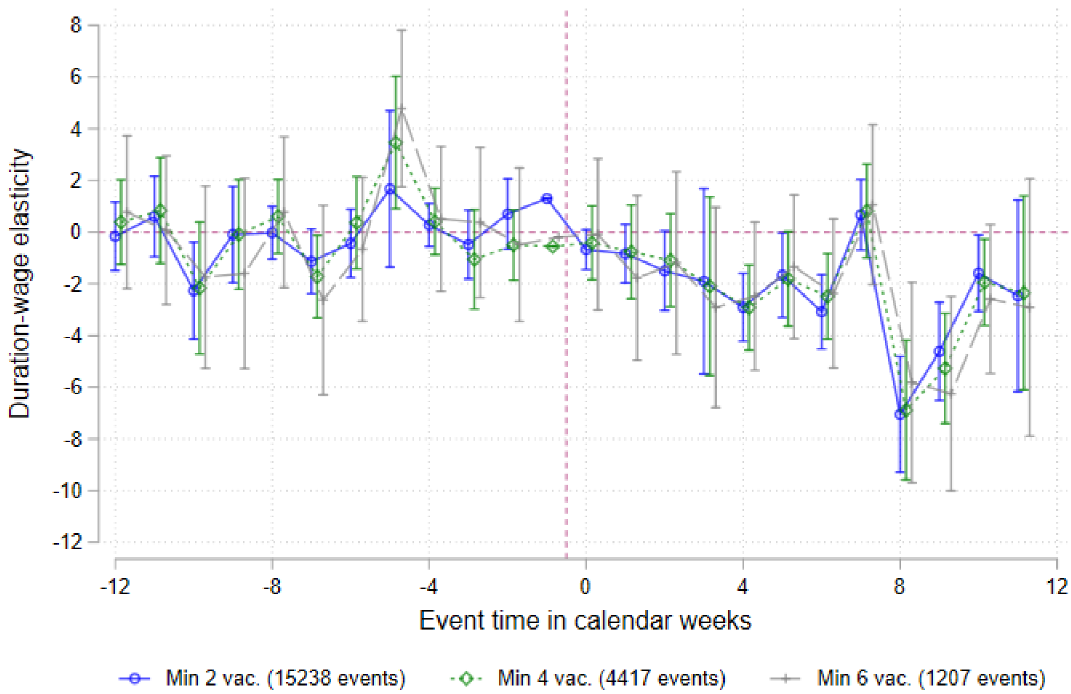

Notes. The figure shows robustness on the event-study coefficients shown in figure 1, by selecting events under alternative criteria on the number of vacancies available in the event window. Imposing a minimum of 2 vacancies (1 pre and 1 post) yields 15,238 wage-change events in our sample, according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

Figure A14: **Event study estimates: Leave-one-out specification**

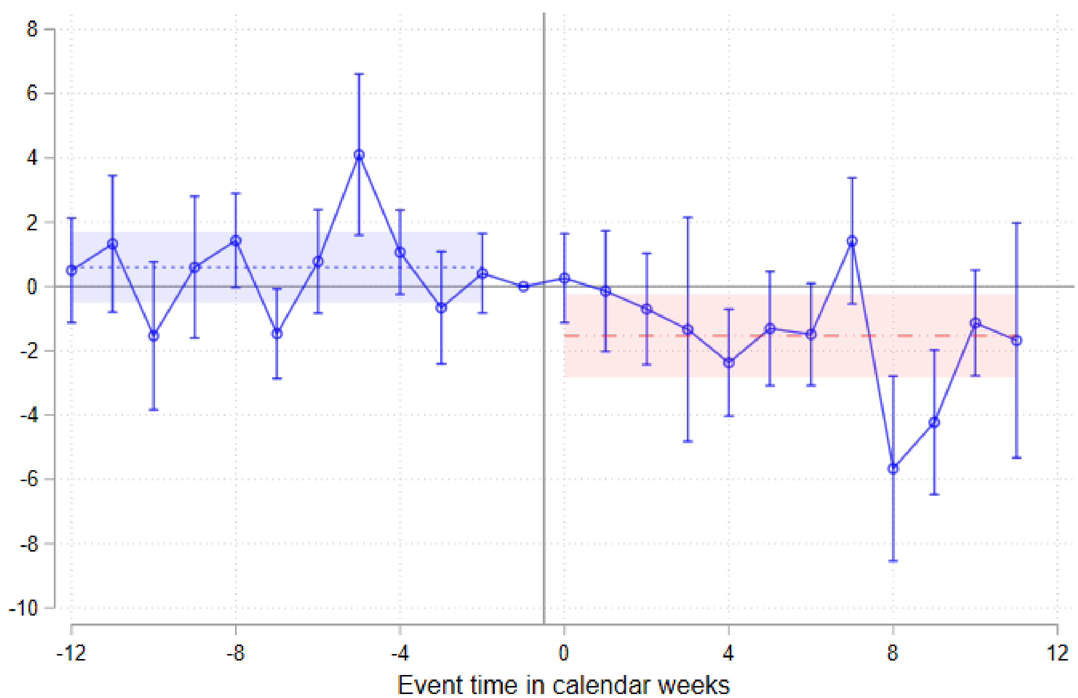

Notes. The figure shows robustness on the event-study results shown in Figure 1, by using as instrument the leave-one-out mean of firm-level wage changes. All other details are as in Figure 1.

Table A1: **Descriptive statistics on estimation samples**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Sample	All Raw	All Clean	Baseline	External Treat	External Control	Internal Treat	Internal Control
Vacancies (th.)	52229	32524	21660	66	333	19	211
Jobs (th.)	24797	16538	6038	19	114	3.5	7.6
Wage (th., mean)	37.9	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)	17.8	18	16.3	13.8	14.8	17.6	11.6
Duration (p50)	15	17	13	11	11	18	8.7
Occupation (pct)							
High skill	47.4	47.6	47.1	25.3	43.4	27.2	42.8
Mid skill	27	27.9	28.7	21.6	27.5	23.7	31.8
Low skill	25.6	24.4	24.3	53.1	29.2	49.1	25.4

Notes. The Table describes different samples of the cleaned Adzuna vacancy-level data for 2017 to 2019. Column 1 refers to all adverts in the raw sample, excluding only those whose recorded duration does not match the number of observed vacancy posts (about 3 million). Column 2 refers to adverts with non-missing wages. Column 3 is restricted to the sample with at least two adverts per job; this sample is used for baseline estimates in Table 1. Column 4 refers to the sample of firms in the Adzuna data that can be matched to the LRD pay-settlement database and column 5 refers to 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 events (a change in firm-level wages of at least 5%, surrounded by 12 weeks on either side of nil wage changes) and 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 location (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 A2: **Estimates based on external pay settlements: Robustness analysis**

	(1)	(2)	(3)	(4)	(5)
First stage		0.660**	0.008	0.393***	0.414***
		(0.265)	(0.005)	(0.062)	(0.062)
Reduced form		-2.737**	-0.086***	-1.700*	-1.960**
		(1.257)	(0.030)	(1.027)	(0.997)
Main equation	-0.142***	-4.145**	-10.849	-4.329	-4.732*
	(0.048)	(1.696)	(7.986)	(2.703)	(2.520)
A-R CI		[.,-0.18]	[.,-2.94]	[-10.00,0.92]	[-10.20,0.16]
F-stat		6.196	2.220	39.792	44.541
Job FE	Y	Y	Y	Y	Y
Location trends	Y	Y	Y	Y	Y
Trimmed	Y	Y	Y	Y	Y
Pay set. IV		Y	Y	Y	Y
No magnitude			Y		
Control				P-score	N-N
Vacancies	65869	65869	65869	445677	128618
Jobs	19354	19354	19354	150107	49867

Notes. The table shows robustness on the elasticity estimates presented in Table 2, based on pay settlements in the LRD database. The sample in columns 1 to 2 includes only treated firms, i.e. matched to pay settlements in the LRD database. Column 1 shows the baseline specification (OLS); column 2 uses the magnitude and timing of the wage event as an instrument for the wage in the current vacancy, and column 3 uses its timing alone. Columns 4 and 5 include matched controls (using wages, benefits and location-specific trends as covariates) and report estimates based on propensity-score weights (column 4) and nearest neighbor matching (column 5). 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 A3: **Estimates based on internally-defined wage events: Wage events defined on at least 10 vacancies**

	(1)	(2)	(3)	(4)
First stage			0.886***	0.693***
			(0.050)	(0.112)
Reduced form			-3.114***	-2.933***
			(0.930)	(0.871)
Main equation	-0.074***	-0.124***	-3.514***	-4.231***
	(0.004)	(0.005)	(1.122)	(1.411)
A-R CI			[-7.73,-0.63]	[.,-0.93]
F-stat			315.940	38.627
Date X TTWA FE	Y	Y	Y	Y
Job FE	Y	Y		
Event FE			Y	Y
Trimmed		Y	Y	Y
Firm wage IV			Y	Y
Leave-one-out				Y
Vacancies	236592	232226	224484	207325
Jobs	12179	12167	10619	9795

Notes. The Table shows estimates on the same specifications shown in Table 3, having defined wage-change events based on a minimum of 10 vacancies (as opposed to 3). All other details are as in Table 3.

Table A4: **Estimates based on internally-defined wage events under alternative control samples**

	(1)	(2)	(3)	(4)
First stage		0.867***	0.778***	0.812***
		(0.048)	(0.116)	(0.077)
Reduced form		-0.912**	-2.172***	-2.057***
		(0.369)	(0.798)	(0.658)
Main equation	0.222**	-1.051**	-2.791***	-2.533***
	(0.107)	(0.434)	(1.063)	(0.808)
A-R CI		[-1.89,-0.13]	[., .]	[.,0.06]
F-stat		320.973	44.685	112.017
Date X TTWA FE	Y	Y	Y	Y
Event FE	Y	Y	Y	Y
Firm wage IV		Y	Y	Y
Control			P-score	N-N
Vacancies	19055	19066	1308437	40246
Jobs	3949	3955	248410	13036

Notes. The table shows robustness on the elasticity estimates presented in Table 3, based on internally-defined wage events. The sample in columns 1 to 2 includes only treated firms, i.e. that experience a large wage change at time t and nil changes in the surrounding 24 weeks. Column 1 shows the baseline specification (OLS) and column 2 uses the magnitude and timing of the wage event as an instrument for the wage in the current vacancy. Columns 4 and 5 include matched controls (using wages, benefits and location-specific trends as covariates) and report estimates based on propensity-score weights (column 4) and nearest neighbor matching (column 5). 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.