I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 13691

Flexible Work Arrangements in Low Wage Jobs: Evidence from Job Vacancy Data

Abi Adams-Prassl
Maria Balgova
Matthias Qian

SEPTEMBER 2020

I Z A Institute
of Labor Economics
Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 13691

Flexible Work Arrangements in Low Wage Jobs: Evidence from Job Vacancy Data

Abi Adams-Prassl
University of Oxford

Maria Balgova
IZA

Matthias Qian
University of Oxford

SEPTEMBER 2020

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Flexible Work Arrangements in Low Wage Jobs: Evidence from Job Vacancy Data*

In this paper, we analyze firm demand for flexible jobs by exploiting the language used to describe work arrangements in job vacancies. We take a supervised machine learning approach to classify the work arrangements described in more than 46 million UK job vacancies. We highlight the existence of very different types of flexibility amongst low and high wage vacancies. Job flexibility at low wages is more likely to be offered alongside a wage-contract that exposes workers to earnings risk, while flexibility at higher wages and in more skilled occupations is more likely to be offered alongside a fixed salary that shields workers from earnings variation. We show that firm demand for flexible work arrangements is partly driven by a desire to reduce labor costs; we find that a large and unexpected change to the minimum wage led to a 7 percentage point increase in the proportion of flexible and non-salaried vacancies at low wages.

JEL Classification: J23, J31, J80

Keywords: flexible jobs, minimum wage, labor demand

Corresponding author:
Maria Balgova
IZA
Schaumburg-Lippe-Straße 5-9
53113 Bonn
Germany

E-mail: balgova@iza.org

* We are very grateful for financial support from UKRI (ES/S010424/1).

1 Introduction

Why do employers offer "flexible" jobs? In many developed economies, the extent to which job flexibility is a positive amenity that reflects workers' demand for work-life balance versus a cost-saving measure used by employers to shift risk onto employees is a matter of fierce policy debate.[1] As described in an independent review of work practises for the UK government: "being able to work when you want is a good thing; not knowing whether you have work from one day to the next when you have bills to pay is not."

Understanding the drivers of flexible work arrangements has been held back by a lack of adequate data and measurement difficulties. The existing empirical literature on job flexibility is largely experimental, focusing on estimating worker preferences for flexibility in hypothetical or tightly controlled scenarios (Mas and Pallais 2017; Wiswall and Zafar 2018; Datta 2019). Administrative data rarely captures contract type or flexibility. Survey questions on work arrangements are not asked consistently over time (Mas and Pallais 2020) and nuanced legal definitions of various employment contracts can make it difficult to interpret survey responses. For example, the UK's Office for National Statistics has admitted that zero-hours contracts (one of the headline forms of flexible work arrangements under which neither employer nor employee commits to providing work) are under-recorded by the Labour Force Survey because workers do not always identify with the precise labels used to describe the work pattern (Adams, Prassl, et al. 2018).

In this paper, we analyze firm demand for flexible jobs by exploiting the language used in job-vacancies. We define a flexible work arrangement as one in which the timing of work is not fixed in the contract and has to be agreed at a later date between the employer and the employee. We consider flexibility alongside whether the advertised job is permanent, full-time, and, crucially, salaried or non-salaried. The implications of flexibility for worker welfare depend on the interaction of these job characteristics. For salaried, full-time jobs, flexibility does not translate into earnings uncertainty for workers, and simply allows more freedom to arrange work around business needs or family commitments. However, in jobs that are paid by the hour, scheduling flexibility can generate uncertainty over earnings.

Working with more than 46 million job vacancies advertised in the UK between 2014 and 2019, we build an information retrieval system in the tradition of Schütze, Manning, and Ragha-

[1] See, for example: the UK review of modern working practises, `https://www.gov.uk/government/publications/good-work-the-taylor-review-of-modern-working-practices`.

van (2008). We take a supervised machine learning approach to classify the work arrangements described in the vacancy text; we manually label the work arrangements described in ∼7,000 job vacancies and use these as the training dataset for a logistic classification model with Lasso regularization, selecting the tuning parameter by a grid search over a 5-fold cross-validated balance F-measure. Our model accurately predicts many dimensions of work arrangements, both in absolute terms, but also relative to the use of hand crafted rules based on keywords.

We apply the trained model to classify the work arrangements and salary-type described in all 46 million vacancies in our dataset. We find considerable heterogeneity in the prevalence of alternative work arrangements across occupations, geography, and the wage distribution. Flexible arrangements are much more likely to be described in low-paid and non-salaried jobs, and in lower-skilled elementary and sales occupations. In 2014, 41% of jobs advertising a wage less than £7 per hour were schedule-flexible, compared to 13% of vacancies with a wage greater than £14 per hour. Flexible vacancies are significantly less likely to be salaried; while 61% of non-flexible vacancies are salaried, only 35% of flexible vacancies are.

We find that the share of flexible vacancies grew significantly between 2014 and 2019. Our results suggest a growing of polarisation in work arrangements across the wage distribution over this period. The growth of non-salaried flexible jobs was fastest in occupations with low wages and which started with a higher initial share of such vacancies, while the opposite was the case for flexible vacancies that were salaried.

To explore the importance of a cost-shifting motive for employer demand for flexibility, we analyze the impact of a large and unexpected increase in the national minimum wage on the work arrangements posted by firms. We exploit the fact that the exposure to the minimum wage increase varied within occupations across different geographical areas. We find that counties and occupations where the increase in labour costs was most binding saw a significant increase in the proportion of flexible-non-salaried vacancies, while the share of permanent, salaried and full-time vacancies decreased. Importantly, the increase in flexible work arrangements was entirely concentrated amongst non-salaried vacancies, leaving the share of schedule-flexibility among salaried vacancies unchanged. We observe limited spillover effects on work arrangements in higher wage brackets; most of the impact of the minimum wage response fell on vacancies offering £9 or less. As a result, there was a significant shift in the distribution of work arrangements across wage levels, giving rise to a polarised picture in which low wage vacancies became disproportionately more likely to be flexible, non-salaried, and without full-time hours or a permanent contract.

In summary, this paper makes three main contributions to the literature. First, we contribute to the growing literature on alternative work arrangements (Mas and Pallais 2020; Datta, Giupponi, and Machin 2019; Prassl 2018; Boeri, Giupponi, Krueger, and Machin 2020). A recent literature estimates worker preferences over flexible arrangements, finding that the average worker is not willing to pay for schedule flexibility and prefers the characteristics of "traditional jobs" (Mas and Pallais 2017; Datta 2019). We complement this work by focusing on firm demand for flexible work arrangements. We highlight the existence of very different types of flexibility amongst low and high wage vacancies. Job flexibility at low wages is more likely to be offered alongside a wage contract. We also observe that firm demand for flexible, non-salaried contracts rises with an increase in the minimum wage, suggesting a cost-shifting motive for flexibility amongst low wage jobs. Flexibility at higher wages and in more skilled occupations is more likely to be offered alongside a fixed salary, and thus workers are shielded from earnings uncertainty.

Second, this is the first paper to use machine learning methods to analyze work arrangements from job vacancy text. Online job vacancy data has been used to analyze firm demand for skill (Hershbein and Kahn 2018; Deming and Kahn 2018; Clemens, Kahn, and Meer 2020) and changes in firm hiring strategies to public policy interventions (Marinescu 2017; Duchini, Simion, and Turrell 2020), and economic shocks (Javorcik, Stapleton, Kett, and O'Kane 2020; Forsythe, Kahn, Lange, and Wiczer 2020). We extract information from job vacancy text to build a consistent time-series of firm demand for flexible work arrangements. Our machine learning approach achieves high accuracy on a range of metrics and outperforms simple keyword search methods to try and classify work arrangements.

Finally, we contribute to the small literature on the impact of minimum wages on non-wage job attributes (Clemens, Kahn, and Meer 2018; Datta, Giupponi, and Machin 2019). Datta, Giupponi, and Machin (2019) were the first to offer evidence of a relationship between rising national minimum wage and the prevalence of flexible work arrangements: they find that the increase in the minimum wage which we also exploit in this paper led to a 5 p.p. increase in the use of zero-hours contracts amongst UK social care workers. We complement this study to analyze the impact across all occupations and wage levels, and for a wider range of flexible work arrangements (only the minority of flexible jobs are zero-hours contracts).[2] Our results show

[2]Adams-Prassl, Boneva, Golin, and Rauh (2020) find that of all employees with variable hours, 8% were on a zero-hours contract, 40% had variable hours subject to the needs of the business but were guaranteed at least some

that employers use flexible work arrangements as a dimension to reduce labor costs, and that the increase in the national minimum wage caused firms to post a higher proportion of flexible and non-salaried jobs with important implications for the welfare effect of the policy.

The rest of this paper proceeds as follows. In Section 2, we describe our data, outline the classification of work arrangements, and describe the machine learning algorithm used to classify vacancies. In Section 3, we chart the variation in flexible work arrangements across occupations, pay, space, and time. In Section 4, we introduce our conceptual framework of firm demand for flexible jobs, and analyse changes in firms' demand for flexibility in response to the 2016 increase in the national minimum wage. Section 5 concludes.

2 Data & Measurement

Employment arrangements vary along multiple dimensions. Heterogeneity in pay, working hours, and the duration of a contract (permanent or temporary) have all been studied extensively. A recent literature has examined the role of flexibility in work arrangements – the ability of employers and/or workers to adjust when and where is work performed (Mas and Pallais 2020). So-called gig or platform work can be seen as an extreme end of this dimension, allowing both workers and employers almost perfect autonomy in deciding when and how much to work (Hall and Krueger 2018).

Flexible work arrangements present many measurement challenges (Abraham and Amaya 2019; Mas and Pallais 2020). While permanency and work hours have clear-cut definitions,[3] there are many ways in which a job can be flexible; a flexible arrangement may refer to flexitime and remote working as well as shift patterns and ad-hoc work. It is known that recording of alternative work arrangements in survey data is sensitive to the wording used and that prompting is often required for accurate reporting (Adams, Prassl, et al. 2018; Abraham and Amaya 2019).

work each week, and 52% had an arrangement in which the worker had the autonomy to determine hours and scheduling.

[3]Duration and work hours have clear-cut defintions: Permanent contracts are defined as contracts without a pre-specified end date. The default for a permanent contract is that the employer-employee relationship continues, in contrast to temporary and fixed-term contracts where the default is a termination. The standard approach is to classify working hours into full-time (30 hours a week or more) versus part-time contract (less than 30 hours a week).

2.1 Burning Glass Technologies Data

To build a consistent time-series on firm demand for flexible work arrangements, we analyze the language used in job-vacancies. Working with an extensive set of job vacancies advertised in the UK, we build an information retrieval system in the tradition of Schütze, Manning, and Raghavan (2008) to classify the work arrangements described in the vacancy text. We use the UK NOVA data feed from Burning Glass Technologies (BGT) as our corpus of job vacancies, a widely used database that aims to cover the universe of online job postings. It is increasingly used to provide real-time information on labour market trends (Javorcik, Stapleton, Kett, and O'Kane 2020; Forsythe, Kahn, Lange, and Wiczer 2020) and to provide new insights into dimensions of firm labour demand (Hershbein and Kahn 2018; Deming and Kahn 2018; Clemens, Kahn, and Meer 2020). Since 2012, BGT has collected the near-universe of electronically posted job vacancies in the UK by scraping 7,500 online job boards and company web pages. We restrict our analysis to data from 2014; between 2012 and 2014, the number of webpages that BGT scraped for job adverts increased rapidly but has been stable since this date. Our analysis is thus based on the full vacancy text of over 46 million unique online job postings between 2014 and 2019.

While the number of vacancies covered by BGT is vast, the data comes with some caveats. The source only captures jobs advertised online. However, some, usually low-skilled jobs, are advertised only in newspapers and magazines. Moreover, other jobs may not be advertised at all; some vacancies might be filled via recommendations within social networks or through head hunters. Smaller firms, whose company web pages are not among those monitored by Burning Glass, may choose to advertise jobs only on their company webpage and not on job boards.[4] Similarly, the data from Burning Glass does not include jobs advertised on closed platforms such as the Civil Service portal or LinkedIn.

Further, there is not necessarily a neat mapping between the prevalence of particular contractual types advertised by firms and the distribution of work arrangements within filled jobs. Not all jobs that are advertised are filled and the terms negotiated in the course of filling a job might deviate from those described in the vacancy itself. We are thus careful to emphasise that our measure reflects firm demand for various contractual arrangements, rather than those that are realised once a match is made. In the future, we hope to match the vacancy data with employer-employee data

[4] One area where Burning Glass Technologies misses a large number of job vacancies posted online is for start-ups. For example, there is not a single job vacancy listed within the corpus of Burning Glass for Eigen Technologies, a LawTech firm, which lists 150 employees on LinkedIn.

to facilitate a full analysis of the distribution of work arrangements in realised matches.

2.2 Classifying Work Arrangements

Our goal is to retrieve all vacancies that describe flexible work arrangements from the set of BGT job adverts. We take a supervised machine learning approach that relies on manual annotations. The annotations serve as a training dataset for our machine learning model and create a "ground truth" for the correct classification of the work arrangements described in job vacancies. Our method proceeds as follows:

1. Manually label a set of job vacancies for the dimensions of work arrangements of interest;
2. Define the vocabulary and represent each job vacancy in a matrix format;
3. Train a machine learning model to classify work arrangements on the basis of vacancy text;
4. Apply the machine learning model to all 46 million job vacancies.

Manual Annotations. We manually label a set of ~7,000 job vacancies with which work arrangements are described in the vacancy text. We use this as a training dataset for a machine learning model to extend the classification to all 46 million vacancies.

We treat flexibility as *schedule flexibility*, i.e. any arrangement in which the timing of work is not fixed in the contract and has to be agreed at a later date between the employer and the employee. While there have been attempts at identifying which party has the power to set the schedule (Adams-Prassl, Boneva, Golin, and Rauh 2020), this is not important for our definition of schedule flexibility. The key aspect is that the work schedule has to be decided on *after* the contract has been signed, leaving it open to future (repeated) negotiations. In practice, we categorize a job to be schedule flexible if it mentions shift or rota work without a fixed pattern, offers flexible working (including remote work), or specifies that work will be organised according to the needs of the business.

It is important to consider schedule-flexibility alongside other characteristics of a work arrangement. In particular, schedule flexible jobs without a fixed salary can give rise to variability in earnings if realised hours vary. A prime example of this is the zero hours contract, but this could also effect workers on part-time contracts where the exact number of working hours fluctuates and

Table 1: Definitions & Examples of Work Arrangements

Dimension	**Description**	**Example**
Schedule flexible	The job has flexible work hours.	"Hours will be weekdays with some flexibility required depending on variable teaching and meeting commitments." "As this is a Bank position to provide cover as and when we need it the hours and days you work will vary." "Ad hoc working- which allows flexibility and choice."
Permanent	The job does not have a pre-specified end date.	"Our global law firm client is currently recruiting for a Senior Server Analyst to join their team on a permanent basis."
Full-time	The job is for more than 30 hours per week.	"The school are looking for a highly experienced and successful primary teacher to join them on a full time basis."
Salaried	The job ad states the payment information as an annual salary.	"Starting salary from £24,300 - £24,500 per annum with further progression opportunities to £24,800."

can be renegotiated. We therefore also label vacancies according to whether they are salaried (i.e. guarantee a fixed income), permanent, and full-time.

Table 1 gives some examples of text describing each of the dimensions we consider. At times, the vacancy text can contain ambiguous or contradictory statements about whether the vacancy is of a particular work arrangement. We address this ambiguity by having two researchers label each vacancy independently who then reconcile any disagreements in their flags in a mutual discussion. We continued to annotate vacancy text until we were satisfied with the accuracy of our information retrieval system. Table 2 gives the number and share of vacancies in our training set where a particular work arrangement was present.

Table 2: Summary Statistics for Manual Vacancy Annotations

Dimension	**Vacancies flagged**	**Vacancies read**	**Share flagged**
Schedule flexible	1710	6376	0.268
Permanent	2953	6650	0.444
Full-time	2803	6679	0.420
Salaried	439	1192	0.368

Recasting Vacancy Text in Matrix Format We treat each dimension of a work arrangement as a separate query for our information retrieval system. Given a query, our goal is the retrieval of all vacancies that mention the specific employment arrangement from the set of BGT job adverts. To apply machine learning classifiers to the training data, we represent the vacancy text in matrix format. To do so, we first define our 'vocabulary', the set of relevant language components. The terms of the vocabulary are determined by the tokenization of the text of our manually annotated vacancies. Tokenization is a way of separating a piece of text into smaller units ('tokens'). We use the tokenisation of 1-grams (single words) to build a parsimonious and interpretable machine learning model. To this, we feature engineer 2-grams and 3-grams on the basis of key sentences highlighted when annotating the training set; that is we hand pick phrases to build a vocabulary which can better differentiate among employment arrangements. We further tune the vocabulary through the use of a stop word list which filters out common words (e.g. and, the, it) from the vocabulary before further processing. Stop words have limited lexical content and their presence adds a lot of noise and little signal to help us distinguish between the dimensions of employment arrangements.

With our vocabulary in hand, we can represent vacancies with the binary term-document incidence matrix. Each row of the incidence matrix is a term of our vocabulary and each column is a vacancy in the training set. Equation 2.1 gives an example of the structure of this matrix. Our incidence matrix is binary as supposed to continuous, as we do not count the number a words that appear in the text, but instead only assign a one if a word appears in the text and a zero if it does not.[5] As the incidence matrix does not take into account the order of the words as they appear in

[5]We experimented with the term frequency inverse document frequency representation of text, however, our validation accuracies were superior using the binary term-document incidence matrix.

the sentence, it resembles a bag-of-words approach. Finally, we limit the size of our vocabulary, and therefore the number of rows of the incidence matrix, to be 5000, to take into account the limited number of annotations, which are as low as 1192 for the payment frequency classes.

$$X = \begin{array}{cl} \left[\begin{array}{cccc} 1 & 0 & \dots & 1 \\ 1 & 1 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & 1 \end{array}\right] & \begin{array}{l} \text{Flexible} \\ \text{Maternity} \\ \text{Fixed Term} \\ \vdots \\ \text{Permanent} \end{array} \end{array} \tag{2.1}$$

Logistic Prediction Model with Lasso Regularization. The incidence matrix corresponds to our matrix of regressors for predicting the work arrangements, where each token is a regressor. As the number of regressors is large, and can be larger than the number of observations, we use a logistic regression model with lasso regularisation. This algorithm chooses a subset of the original 5000-token vocabulary to predict work arrangements from the vacancy text to avoid over-fiiting.

Let $\mathbb{C}$ be the set of C classes that we want to predict (in our case the different employment arrangements) and $\mathbb{D}$ represent the set of D vacancies in the training set. Let $c_i = 1$ if vacancy i has been labelled with work arrangement $c \in \mathbb{C}$ and $c_i = 0$ otherwise. x_i represents the vector representation of document $i \in \mathbb{D}$ using our vocabulary. With the logistic specification, the probability of a vacancy i being labelled with class c given the representation of its text, x_i, is:

$$\Pr_c(c_i = 1|x_i) = \frac{\exp(x_i'\beta_c)}{1 + \exp(x_i'\beta_c)}. \tag{2.2}$$

Under lasso regularization, the log-likelihood of the logistic regression model is:

$$\sum_{i=1}^{D} \log \Pr(c_i|x_i) - \frac{\lambda}{2}||\beta_c||_1 \tag{2.3}$$

where λ is the tuning parameter which governs the magnitude of the penalty term. Intuitively, with sufficiently high λ, the lasso regularization penalizes the inclusion of non-zero coefficients on vocabulary elements.

Algorithm 1 lists the steps taken to train the logistic regression model. In summary, we determine the tuning parameter, and thus model parameterization, by a grid search over a 5-fold cross-validated balance F-measure. Consider a work arrangement c. For a given value of the tuning parameter λ, we randomly sample a balanced set of vacancies (i.e. containing an equal number of zeros and ones for that class) from the training set. Call this sample $\mathbb{D}_S$. We split $\mathbb{D}_S$ into five equally sized groups $\mathcal{S}_j$ for $j = 1, ..., 5$. For each group $\mathcal{S}_\ell$, we select the parameter, $\hat{\beta}_{c,\lambda,\ell}$ that maximises the objective function at Equation 2.3 on the subset of vacancies excluding group $\mathcal{S}_\ell$, i.e. $\mathbb{D}_S \setminus \mathcal{S}_\ell$. $\hat{\beta}_{c,\lambda,\ell}$ is then used to predict the presence of work vacancy c in the excluded group $\mathcal{S}_\ell$. The model's predictions are compared to the true vacancy labels for vacancies in $\mathcal{S}_\ell$ and the accuracy of the model evaluated (see below). We repeat this procedure 500 times for each value of λ for each work arrangement. We select the value of λ that maximises the average of the cross-validated balanced F-measure over the 500 sampling draws.

We use the balanced F-measure to evaluate the accuracy of our classification algorithm. This equally weights precision and recall using the harmonic mean:

$$\text{Balanced F-measure} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}. \tag{2.4}$$

where

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \tag{2.5}$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{2.6}$$

and TP, FP, and FN are defined according to the congruence of a vacancy label and the model prediction as follows:

		True Label	
		0	1
Predicted Label	0	True Negative (TN)	False Negative (FN)
	1	False Positive (FP)	True Positive (TP)

The precision score thus penalises false positives (irrelevant items retrieved) and the recall score penalises false negatives (relevant items that are not retrieved).

Table 3 shows the precision, recall and balanced F-measure for our trained model. We have a consistently strong performance across all dimensions of work arrangements, both in absolute terms, but also relative to the use of hand crafted rules based on keywords. The keywords used in the benchmark correspond to the most positive keyword identified by the logistic regression model (displayed in Table A.1). The outperformance of our machine learning model is due to the positive and negative word weights in the logistic regression model, while the simple hand crafted rules only use the positive word weights.

Algorithm 2 lists the steps to apply our trained machine learning model to predict the employment arrangement given a previously unseen vacancy text. The trained machine learning model is used to evaluate the probability that a work arrangement is described in a given vacancy, $\Pr_c(1|x_i)$. If this probability is greater than .5, the vacancy is labelled with that particular work arrangement.

Algorithm 1: Training the Logistic Regression model $(\mathbb{C}, \mathbb{D})$

Result: Return the tuple $W = (\hat{\beta}_1^*, \hat{\beta}_2^*, ..., \hat{\beta}_C^*)$ containing the vocabulary weights for each class $c \in \mathbb{C}$

extract the vocabulary in $\mathbb{D}$

for $c \in \mathbb{C}$ **do**

  for $\lambda \in [0.01, 0.05]$ **do**

    let $j = 0$

    repeat

      compute as N_c the number of vacancies in class c which are labelled with 1

      sample N_c vacancy texts labelled with 0 and the same number of text labelled with 1 and call the sample $\mathbb{D}_S$

      randomly divide documents in $\mathbb{D}_S$ into 5 equally sized blocks $\mathcal{S}_j$, for $j = 1, ..., 5$

      for $j = 1, ..., 5$, with $\Pr_c(c_i|x_i)$ defined as in Equation (2.2), find $\hat{\beta}_{c,\lambda,\ell}$ which minimises

$$\sum_{i \in \mathbb{D}_S \setminus \mathcal{S}_\ell} \log \Pr_c(c_i|x_i) - \frac{\lambda}{2} ||\beta_{c,\lambda}||_1$$

      where $||\cdot||_1$ is the 1-norm while excluding the observations in block ℓ

      compute the cross-validated balanced F-measure, as defined in Equation (2.4), using observations in block ℓ and $\hat{\beta}_{c,\lambda,\ell}$

      increase j by 1

    until $j = \mathit{500}$;

  end

  select the λ with the lowest average cross-validated balanced F-measure and call the corresponding estimate $\hat{\beta}_c^*$

end

Table 3: Prediction Accuracy of Trained Logistic Regression Model

	Logistic Regression Model			Improvement to keywords		
Contract type	**Precision**	**Recall**	**F**	**Precision**	**Recall**	**F**
Schedule flexible	0.8540	0.8083	0.8303	0.0005	0.4399	0.1855
Permanent	0.9294	0.9736	0.9510	0.0471	-0.0067	0.0223
Full-time	0.9162	0.8881	0.9019	0.1898	0.2236	0.1314
Salaried	0.8604	0.8415	0.8503	-0.1032	0.3586	0.2070

Algorithm 2: Apply the Logistic Regression model $(\mathbb{C}, V, W, d)$

Result: Assign class $c \in \mathbb{C}$ to vacancy i

construct a vector representation x_i according to the vocabulary

for $c \in \mathbb{C}$ **do**

compute the probability $\Pr_c(0|x_i)$ and $\Pr_c(1|x_i)$ using the coefficients $\hat{\beta}_c^*$ in weight vector W as

$$\Pr_c(1|x_i) = \frac{\exp(x_i'\hat{\beta}_c^*)}{1 + \exp(x_i'\hat{\beta}_c^*)}, \quad \Pr_c(0|x_i) = 1 - \Pr_c(1|x_i)$$

if $\Pr_c(1|x_i) > \Pr_c(0|x_i)$ **then**

assign class c to vacancy i

end

end

3 Variation in Flexible Work Arrangements

We apply our trained machine learning model to the 46 million unique job vacancies in the BGT dataset posted in the UK between 2014 and 2019. We first examine the cross-sectional distribution of flexible vacancies in our base year (2014) before turning to examine changes in the proportion of flexible vacancies over time.

3.1 Cross-sectional variation

Our dataset captures 5.2 million vacancies in 2014, 18% of which were described as having a flexible schedule. There is considerable variation in the proportion of flexible vacancies over space, occupations, and wages. Turning first to spatial variation, our primary geographic unit is the county; there are 100 counties in the UK, with an average of 696,488 inhabitants each (488,717 at the median) in 2014. Figure 1 shows the distribution of schedule-flexible vacancies across UK counties. The proportion of flexible vacancies advertised in 2014 ranges from 5% in the City of London, England to 43% in the Orkney Islands, Scotland.

Figure 1: Distribution of Flexible Vacancies by County

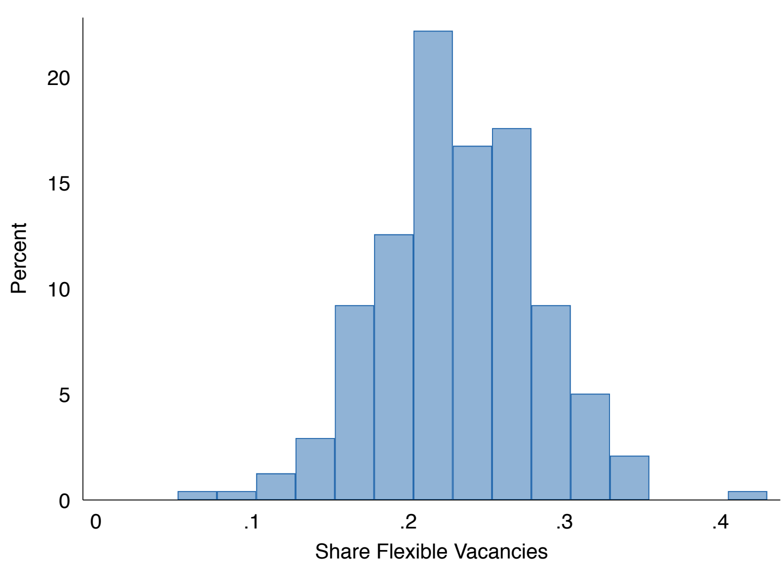

Notes: Figure shows a histogram of the share of schedule-flexible vacancies of all jobs advertised in a county in 2014.

There is a lot of variation in the extent to which flexible scheduling is advertised across different occupations and wage levels. In Figure 2 we report the share of flexible vacancies in 2014 by one-digit SOC codes and wage levels. The share of vacancies reporting flexible scheduling ranges from 44% for 'Caring, Leisure & Other Service Occupations' to 12% for 'Managers, Directors and Senior Officials. The proportion of flexible vacancies in adverts that cannot be assigned an occupation code (18.1%) is in line with the average across those that can (18.4%).

Flexible vacancies are significantly more likely to be advertised in low-wage jobs: 41% of jobs advertising a wage less than £7 per hour were schedule-flexible, compared to 13% of vacancies with a wage greater than £14 per hour. 35% of vacancies were missing wage information in 2014. Vacancies without wage information are significantly less likely to describe a flexible work-arrangement on average; 18% of vacancies missing wage information were flexible, which is more comparable to the level observed in high-wage than low-wage jobs.

Figure 2: Share of Flexible Vacancies by Occupation and Wage

(a) Occupation

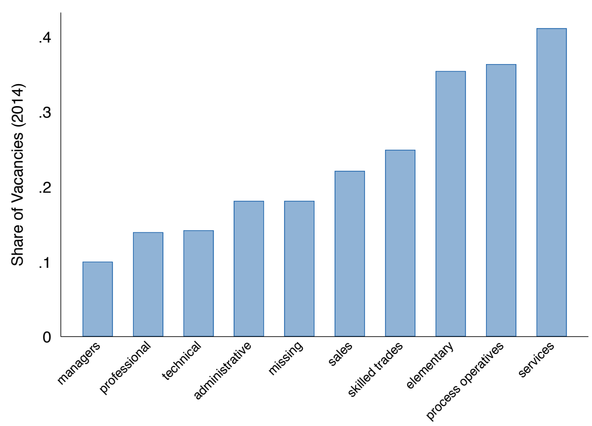

(b) Wage (£per hour)

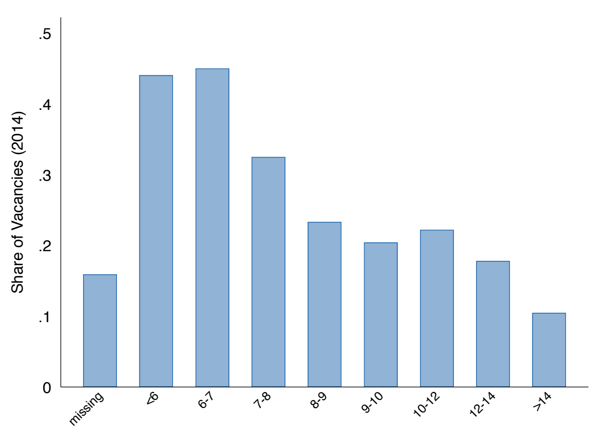

Notes: The bars show the proportion of flexible vacancies by one-digit SOC group and wage group in 2014.

3.2 Flexibility & Remuneration Type

With risk averse workers, it is not just the wage level but also variation in earnings that matters for welfare. The relationship between schedule flexibility and variation in earnings is mediated by whether a job is salaried or non-salaried. When work is paid by the hour, there is more scope for variation in hours and schedules to translate into changes in earnings.

To understand the interaction between pay and flexibility, we restrict our sample to vacancies with wage information. Amongst these vacancies, salaried positions are much more frequent that non-salaried ones. However, this does not extend to flexible jobs, which are less less likely to be advertised with an annual salary compared to non-flexible jobs: while 61% of vacancies with wage-information are advertised as salaried positions, this is true for only 35% of schedule-flexible vacancies. In Figure 3, we further break down the share of salaried and non-salaried schedule flexible vacancies by occupation and wage; the red bars give the proportion of salaried flexible vacancies and the blue bars give the proportion of non-salaried flexible vacancies. Flexible vacancies are less likely to be salaried at low wages and in lower occupation codes: 17% of flexible vacancies advertised at less than £7 per hour were salaried compared to 51% with a wage greater than £14 per hour.

Figure 4 shows the relationship between the share of flexible vacancies and the share of permanent and low wage vacancies at the 3-digit SOC code. Panel (a) confirms the pattern that flexible & non-salaried vacancies are more likely to be advertised in low-wage occupations. In panel (b), we see that fine-grained occupations that post a high share of flexible & non-salaried vacancies are less likely to advertise permanent positions. However, there is no systematic relationship between the share of flexible & salaried vacancies with the share of permanent or low-wage vacancies at the 3-digit SOC code level. These patterns are in line with our hypothesis that the interaction between the remuneration type and schedule flexibility results in very different types of vacancies.

In Table 4 we analyze the variation in the proportion of flexible vacancies across wages and occupations in a linear probability model framework. Column (1) confirms that, even after controlling for county fixed effects, flexible vacancies are substantially more likely to be advertised at low wages (compared to vacancies advertised with remuneration above £14 per hour), without a salary, and in service and process-operative occupations. The next two columns highlight the difference between flexible & salaried vacancies (column (2)) and flexible & non-salaried vacancies (column (3)). There is very little variation in the share of flexible & salaried vacancies (as a proportion of all vacancies advertised) across different occupations and wage bins, in contrast to flexible & non-salaried vacancies which exhibit a strong decreasing relationship with wage, and a relatively larger variation across occupations. However, this difference may be driven by the patterns in salaried vacancies, rather than due to the differences in flexible vacancies. To test this, we repeat the regressions for salaried and non-salaried vacancies separately. Amongst salaried vacancies (column (4)), schedule-flexible jobs are more likely to be advertised at hourly wages lower than £7 and in services & process-operative occupations, although this variation is still of a smaller magnitude when compared to what is observed within non-salaried vacancies (column (5)).

Figure 3: Share of Flexible Vacancies by Occupation and Wage

(a) Occupation

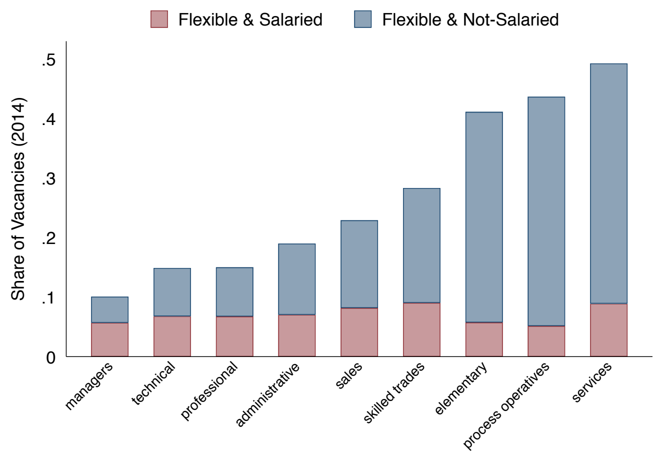

(b) Wage

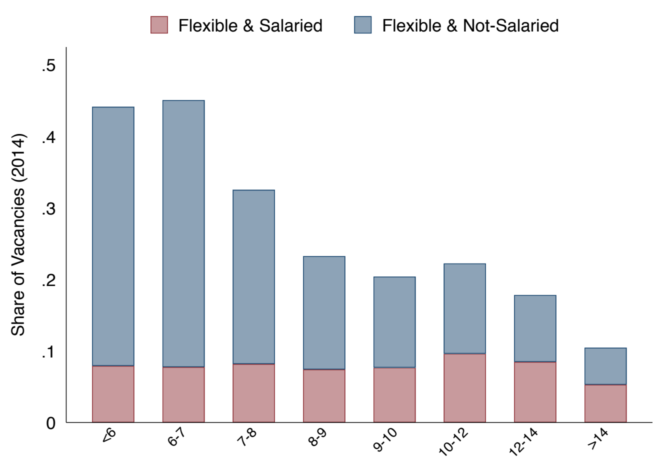

Notes: The bars show the proportion of flexible vacancies by one-digit SOC group and wage group in 2014. The red bar gives the proportion of flexible and salaried vacancies, while the blue bar gives the proportion of flexible and non-salaried vacancies.

Figure 4: Share of Flexible Vacancies by 3-Digit Occupation

(a) Relationship with Low-Wage Vacancies

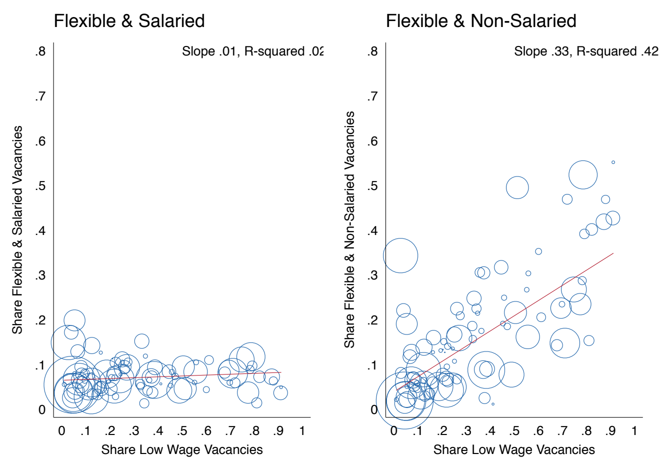

(b) Relationship with Permanent Vacancies

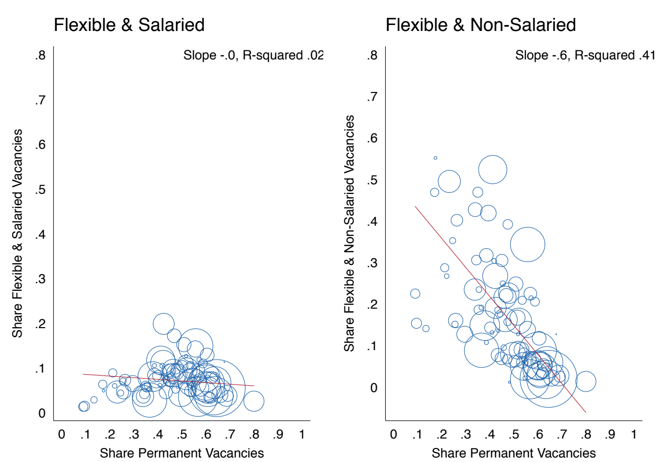

Notes: Each circle represents a 3-digit SOC code with the size proportional to the number of vacancies in 2014 in that occupation. The red line gives the line of best fit. The share of low wage vacancies is the proportion of vacancies where the maximum salary is less than £9 per hour.

Table 4: Characteristics of Schedule Flexible Vacancies

	All			Salaried	Non-Salaried
	Flexible	Flexible & Salaried	Flexible & Non-Salaried		
	(1)	(2)	(3)	(4)	(5)
Wage Level (£; Base = > 14)					
<6	0.2710***	0.0198***	0.2512***	0.2094***	0.2300***
	(0.0014)	(0.0008)	(0.0013)	(0.0025)	(0.0017)
6-7	0.2493***	0.0216***	0.2277***	0.1158***	0.2845***
	(0.0012)	(0.0007)	(0.0011)	(0.0016)	(0.0016)
7-8	0.1509***	0.0250***	0.1258***	0.0600***	0.2249***
	(0.0010)	(0.0006)	(0.0009)	(0.0011)	(0.0016)
8-9	0.0832***	0.0182***	0.0650***	0.0227***	0.1831***
	(0.0009)	(0.0006)	(0.0008)	(0.0009)	(0.0018)
9-10	0.0623***	0.0211***	0.0412***	0.0273***	0.1257***
	(0.0009)	(0.0006)	(0.0007)	(0.0009)	(0.0017)
10-12	0.0819***	0.0395***	0.0424***	0.0565***	0.1274***
	(0.0008)	(0.0006)	(0.0006)	(0.0009)	(0.0015)
12-14	0.0523***	0.0280***	0.0242***	0.0319***	0.1092***
	(0.0007)	(0.0005)	(0.0005)	(0.0007)	(0.0014)
Occupation (Base = Managers)					
professional	0.0459***	0.0081***	0.0378***	0.0303***	0.0395***
	(0.0006)	(0.0004)	(0.0004)	(0.0006)	(0.0014)
technical	0.0148***	0.0038***	0.0110***	0.0126***	0.0083***
	(0.0007)	(0.0005)	(0.0005)	(0.0007)	(0.0016)
administrative	-0.0049***	-0.0010	-0.0039***	0.0127***	-0.0404***
	(0.0009)	(0.0006)	(0.0007)	(0.0010)	(0.0018)
skilled trades	0.1025***	0.0163***	0.0862***	0.0694***	0.0809***
	(0.0012)	(0.0008)	(0.0010)	(0.0013)	(0.0021)
services	0.2406***	0.0142***	0.2264***	0.1215***	0.2232***
	(0.0014)	(0.0009)	(0.0013)	(0.0019)	(0.0021)
sales	0.0389***	0.0121***	0.0267***	0.0132***	0.1067***
	(0.0009)	(0.0006)	(0.0007)	(0.0009)	(0.0020)
process operatives	0.2405***	-0.0240***	0.2645***	0.0940***	0.1836***
	(0.0017)	(0.0008)	(0.0016)	(0.0026)	(0.0023)
elementary	0.1435***	-0.0170***	0.1605***	0.0781***	0.0858***
	(0.0016)	(0.0009)	(0.0015)	(0.0024)	(0.0023)
Constant	0.0682***	0.0481***	0.0201***	0.0649***	0.0904***
	(0.0027)	(0.0018)	(0.0022)	(0.0030)	(0.0046)
Observations	3,326,759	3,326,759	3,326,759	2,032,971	1,346,384
R^2	0.1100	0.0080	0.1344	0.0349	0.1227
County F.E.	yes	yes	yes	yes	yes

Notes: Standard errors in parenthesis. Sample restricted to vacancies in 2014 with wage information. Dependent variable in columns (1)-(3) are dummy variables for whether a vacancy is schedule-flexible, schedule-flexible & salaried, schedule-flexible & non-salaried respectively. Dependent variable in columns (4) and (5) are whether a vacancy is schedule-flexible in the sample of salaried and non-salaried vacancies respectively.

3.3 Changes over time

We turn now to changes in the proportion of flexible jobs over time. Between 2014 and 2019, there was a considerable increase in the prevalence of flexible work arrangements in online jobs postings. The share of flexible vacancies grew by 9 percentage points both amongst all vacancies (from 18% to 27%) and for vacancies with wage information (from 20% to 29%). Given the identical trends, and that fact that our focus is on low-wages and the differences between salaried and non-salaried positions, we restrict the sample to vacancies posting wage information for the rest of this analysis.

Figure 5 (a) shows the share of schedule-flexible vacancies over time. The share of both salaried and non-salaried flexible vacancies increased over this period, but, as panel (b) in the figure shows, non-salaried flexible vacancies grew faster than the salaried flexible ones (from 13% to 19% for the former compared to 7% to 10% for the latter).

This growth did not occur uniformly across different occupations and wage levels. Figure 6 shows the change in the share of flexible vacancies against the initial share and the share of low-pay jobs in 2014. The four panels show a markedly different trend for the salaried and non-salaried flexible jobs. The growth in non-salaried jobs was somewhat faster in occupations which already posted relatively high shares of non-salaried flexible jobs in 2014, and in occupations with a high initial share of low-pay jobs. The opposite was true for salaried flexible jobs, which grew predominantly in occupations with relatively higher wages and with low initial shares of flexible jobs. This suggests a degree of polarisation in flexible work arrangements in the UK over this period.

Figure 5: Growth in Share of Flexible and Permanent Vacancies Over Time

(a) Change in Share of All Vacancies

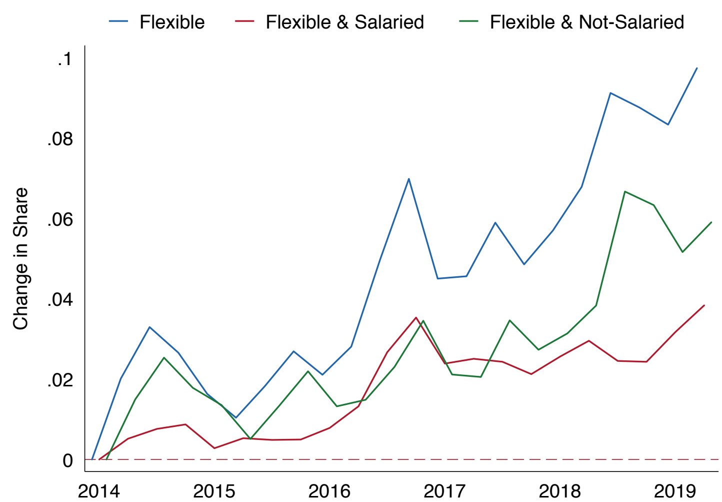

(b) By Salaried & Non-Salaried Vacancies

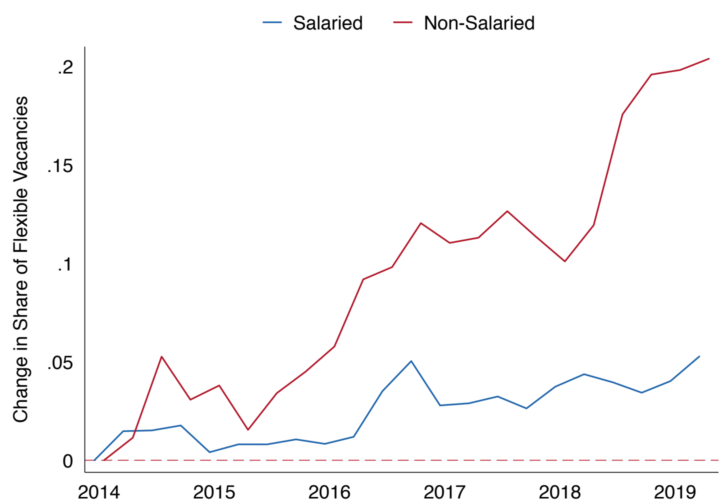

Notes: Line shows the growth in the share of flexible and permanent vacancies in each quarter relative to the first quarter of 2014.

Figure 6: Growth in Share of Flexible Vacancies

By 2014 Share of Flexible Vacancies

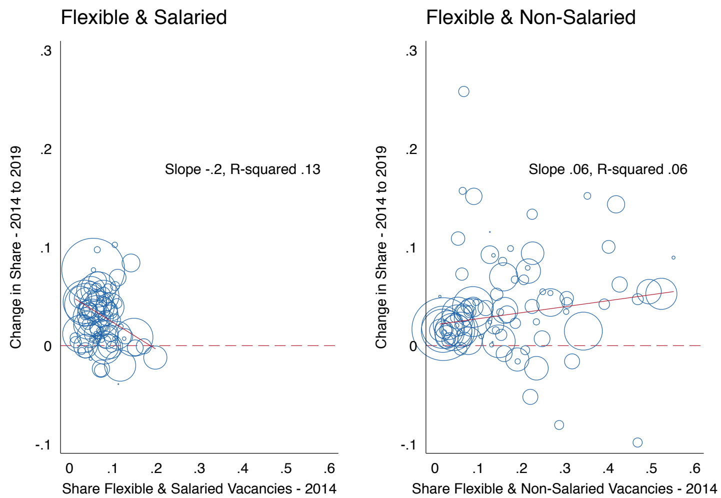

By 2014 Share of Low Wage Vacancies

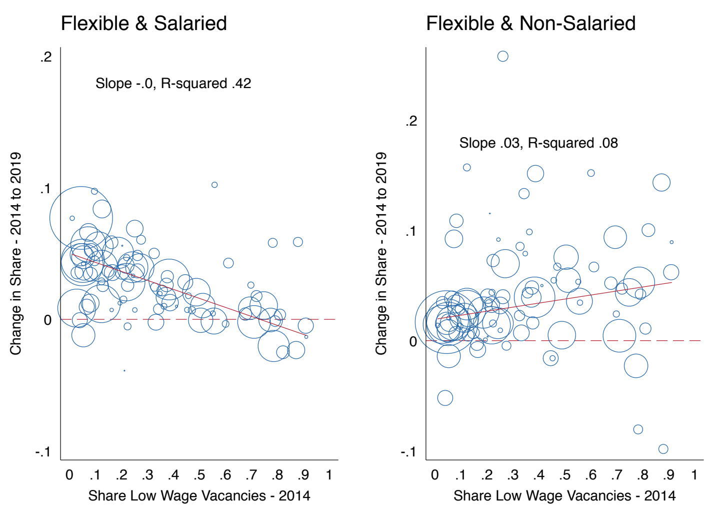

Notes: Each circle represents a 3-digit SOC code with the size proportional to the number of vacancies in 2014 in that occupation. The red line gives the line of best fit. The y-axis gives the change in the share of vacancies between 2014 and 2019.

4 Flexible Jobs & Minimum Wages

Section 3 highlighted that flexible vacancies are concentrated at low wages in non-salaried positions and that there prevalence has been rising over time. A key policy debate in the UK is the extent to which this rise reflects greater worker-led demand for work-life balance versus a new dimension to cut labour costs by employers (Taylor, Marsh, Nicol, and Broadbent 2017). To explore the relationship between firm demand flexibility and labour costs, we analyse the impact of a large and unexpected increase in the national minimum wage in 2016 on the prevalence of these work arrangements.

We proceed in three steps. First, we develop a simple model of firm demand for flexibility in the minimum wage labour market to structure our empirical analysis. Second, we adopt the method of Cengiz, Dube, Lindner, and Zipperer (2019) to confirm that the increase in the minimum wage indeed led to a shift in the distribution of offered wages in job vacancies. Finally, we exploit the regional variation in the bite of the minimum wage to estimate the casual impact of an increase in labour cost on the characteristics of offered work arrangements.

4.1 Conceptual Framework

Let firms face two options when deciding to post a vacancy. The firm can either offer flexible, non-salaried positions that do not commit them to providing a fixed earnings stream to the worker; these effectively mimic the spot-market. Alternatively, firms can hire via traditional fixed-hours (non-flexible) jobs that commit them to provide the worker with a given earnings stream.

This distinction matters because at the time of posting a vacancy, firms face two types of uncertainty: over the level of product demand, and over their ability to hire sufficient labour on the spot market next period. A traditional fixed-hours contract allows the firm to insure against fluctuations in the labour market, but at the cost of having to pay for committed labour when product demand is low. A flexible contract offers the opposite: the option to only pay for the labour the employer needs, at the risk of not being able to secure sufficient flexible workers at short notice when demand is high. What work arrangement is offered depends on which of these two trade-offs is optimal for the employer.

Let product demand be "high" with probability p, or "low" with probability $1 - p$. When product demand is high, the firm is able to sell up to y_H units of its goods; to meet this demand it has to hire l_H units of labour. When product demand is low, the firm only sells y_L goods, and only needs to employ l_L of labour with $l_L < l_H$. If the firm employs labour via flexible contracts, it can satisfy high product demand only with probability r (as with probability $1-r$ workers on a flexible contract will choose not to accept the offer of work). When labour supply is low, the firm can only produce y_L output regardless of the level of product demand. The firm is a price-taker in both the product and labour market, and firm-specific shocks to demand and labour-supply leave both wage and product price unchanged. We normalise product price to 1. All jobs are paid a set wage w, which stands for minimum wage.

By hiring on flexible contracts, the firm only ever needs to pay for the labour it requires. On the other hand, it risks not being able to take advantage of a surge in product demand in the high state. The expected profit of a flexible position is:

$$\begin{aligned} E[\pi|\text{flexible}] &= pr(y_H - wl_H) + (1-pr)(y_L - wl_L) \\ &= pr\pi_H + (1-pr)\pi_L \end{aligned} \tag{4.1}$$

The firm can insure itself against flexible labour turning down work by offering a fixed employment contract for l_H hours. In this case, the firm will receive and pay for l_H labour in all states of the world. The firm earns high profit π_H whenever product demand is high, but it faces a loss relative to the flexible contract when demand is low. The expected profit of a fixed position is:

$$\begin{aligned} E[\pi|\text{non-flexible}] &= p(y_H - wl_H) + (1-p)(y_L - wl_H) \\ &< p\pi_H + (1-p)\pi_L \end{aligned} \tag{4.2}$$

Implications The firm will hire on fixed contracts if the expected forgone profits when product demand is high are greater than the expected overpayment on the wage bill when demand is low. Equations 4.1 and 4.2 imply that inflexible contracts are preferred when:

$$(1-r)p(\pi_H - \pi_L) \geq (1-p)w(l_h - l_L) \tag{4.3}$$

This simple model implies the following. First, firms are more likely to make use of flexible contracts when the cost of labour w increases because it becomes more costly for the employer to hoard labour in times of low product demand. Thus, other things equal, one would expect a rise in flexible non-salaried vacancies with an increase in the national minimum wage.

Second, the prevalence of flexible contracts will be higher when the spot market is relatively thick; the relative advantage of a fixed hours contract is less when it is easy for the employer to hire sufficient flexible labour. Thus, geographic areas and sectors where flexible non-salaried contracts are already relatively widespread will see a relatively larger increase in flexible contracts when such contracts become more attractive, potentially giving rise to a polarisation in the labour market.

Finally, firms are more likely to use flexible contracts with higher variance in product demand. When the probability of facing a large increase in product demand is low, the need to insure against insufficient labour supply is lower.

4.2 Impact on Offered Wages

A national minimum wage (NMW) was first introduced in the UK in 1999. All types of employees are subject to the NMW wage floor; only self-employed workers are outside the scope of the legislation. Table 5 gives the evolution of the NMW since 2014.[6]. The level varies according to the age of the worker, with youth rates set at 76% of the adult rate on average between 2014 and 2019.

We focus on the large and unexpected increase in the NMW in April 2016, when the adult rate was raised from £6.70 to £7.20. This represented a substantial (7.5%) increase in the wage floor for adults. The policy change was announced in July 2015 and was both unexpected and unexpectedly large (Datta, Giupponi, and Machin 2019). The previous literature has found that there is little evidence that this increase led to adverse employment effects nor firm exit (Giupponi and Machin 2018). However, Datta, Giupponi, and Machin (2019) find evidence that it led to a 5 p.p. increase in the prevalence of zero-hours contracts in the social care sector. We analyze the impact of the change on firm demand for a wider range of flexible work arrangements, across all sectors.

We start by confirming that the 2016 increase in the minimum wage lead to a rightward shift in the distribution of posted wages. As wages below £7.20 became illegal for adult workers, we should observe a drop in the share of vacancies advertising a wage below £7, and an increase in the share of vacancies offering more than £7. We focus on the shares of vacancies at different wage levels rather than their absolute number given our interest in the distribution rather than

[6] See `https://www.gov.uk/national-minimum-wage-rates`

Table 5: National Minimum Wage Rates

	Age				
Year	25 and over	21 to 24	18 to 20	Under 18	Apprentice
April 2019 to March 2020	£8.21	£7.70	£6.15	£4.35	£3.90
April 2018 to March 2019	£7.83	£7.38	£5.90	£4.20	£3.70
April 2017 to March 2018	£7.50	£7.05	£5.60	£4.05	£3.50
October 2016 to March 2017	£7.20	£6.95	£5.55	£4.00	£3.40
April 2016 to September 2016	£7.20	£6.70	£5.30	£3.87	£3.30
	21 and over		18 to 20	Under 18	Apprentice
2015	£6.70		£5.30	£3.87	£3.30
2014	£6.50		£5.13	£3.79	£2.73

the overall level of labour demand.

We discretize wage levels into eight bins, with a higher granularity at the lower end of the wage distribution to allow us to capture the effect of the NMW more precisely. For each wage bin, within each occupation, we compare the share of vacancies paid at a given wage level before and after the NMW introduction. We control for time-, season- and occupation- fixed effects, as well as for policy anticipation effects and past minimum wage increases. Following the methodology introduced by Cengiz, Dube, Lindner, and Zipperer (2019), we estimate the following specification:

$$\frac{V_{myow}}{\sum_w V_{myo}} = \sum_w \delta_{ow} After_{my} * I_w + \alpha_{wo} + \beta_{ywo} + \gamma_{mwo} + \omega_{mywo} + \epsilon_{mywo} \tag{4.4}$$

where V corresponds to the number of vacancies in occupation o, month m, year y and wage bin w. α_{wo} is the wage-occupation constant. β_{ywo} and γ_{mwo} are year- and month- fixed effects for each wage-occupation bin. ω_{mywo} are controls for pre-2016 increases in the minimum wage and for the 8 months between the announcement and implementation of the NMW, and ϵ_{mywo} is the error term. The treatment effect of the NMW change on each wage bin is captured by the set of occupation-specific coefficients δ_{ow}. We report robust standard errors clustered at occupation level.

The estimated treatment effects δ_{ow} and their 95% confidence intervals are plotted in Figure 7. Some occupations are much more affected by the minimum wage increase than others. The first three occupation groups (mangers, professionals, and technical) do not register any statistically significant shifts in the posted wage distribution. This is to be expected give that only 2.5% of vacancies in these wage bins were paid below £7.20 per hour in 2015. In the other six occupation groups, the magnitude of the wage distribution shift varies; it is highest in elementary and services occupations, where much of the low-pay vacancies concentrate. Furthermore, the plots suggest that the rightward shift in posted wages is concentrated at just below and just above the NMW thresholds; there is little evidence of spillover effect into wage brackets beyond £9. The first stage thus suggests that the NMW increase represented a significant rise in labour costs for jobs at the lowest end of the wage distribution.

Figure 7: NMW impact on posted wages, by occupation group

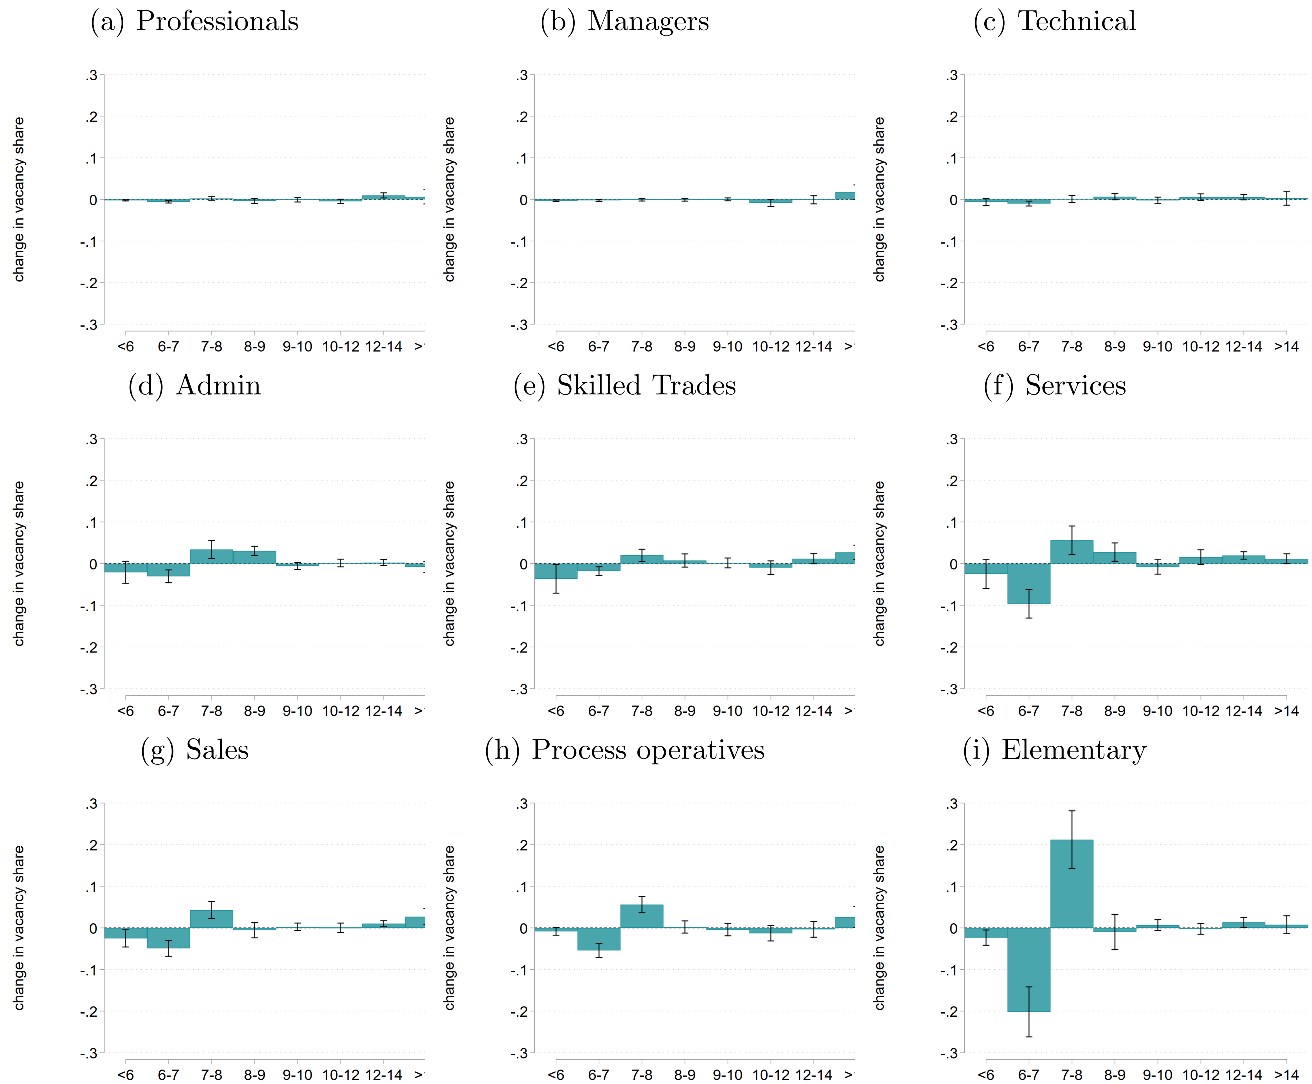

Notes: Figure gives estimates of δ_{ow} from Equation 4.4: estimates of the treatment effect of NMW increase on wage bins within each occupation group. The estimates are identified by comparing the share of vacancies in a given wage bin before and after the NMW introduction, controlling for occupation, season and time FE for each wage bin, as well as for past increases in minimum wage.

4.3 Impact on Flexible Work Arrangements

To establish the impact of the wage shock on advertised work arrangements, we exploit variation in the 'bite' of the NMW change across geographic areas. Our identification strategy is based on the idea that once we control for region fixed effects and occupation-specific time shocks, hiring within a given occupation in high-wage regions is a good control for hiring in the same occupation category in a low-wage region. To improve the comparability of occupation-specific hiring, we use 3-digit SOC occupation categories.

The intensity of treatment, the NMW bite, is calculated as the average share of vacancies paying below £7 in a given occupation-county cell in the year 2014:

$$NMWbite_{oc} = \frac{\text{vacancies paying} < \text{£}7/\text{hour}_{oc}}{\text{all vacancies}_{oc}} \tag{4.5}$$

Figure 8 shows the variation in the bite of the NMW change, aggregated to 1-digit occupation groups. While the NMW bite is close to 0 in all counties for managerial and technical occupations, it is much higher among occupations that are more exposed to the minimum wage increase, and hence from where most of the identifying variation comes from. This variation is similar to the considerable spatial heterogeneity in work arrangements we described in Section 3.1.

Our main specification is:

$$\frac{Y_{ocyq}}{V_{ocyq}} = \delta NMWbite_{oc} * After_{yq} + \alpha_c + \beta_{oyq} + \epsilon_{ocyq} \tag{4.6}$$

The dependent variable $\frac{Y_{ocyq}}{V_{ocyq}}$ captures the proportion of vacancies in an occupation-county-year-quarter cell that describe offering a work arrangement $Y \in \{\text{flexible, salaried, permanent, full-time}\}$. δ is the estimated average treatment effect on the treated for a particular work arrangement Y. It is identified from the interaction between the measure of treatment intensity $NMWbite_{oc}$ and temporal variable $After_{yq}$. $After$ is equal to 1 for second quarter of 2016 onwards, and 0 otherwise. We control for county fixed effects (α_c) and occupation-specific time shocks (β_{oyq}). Standard errors are clustered at occupation-year-quarter cell.

Figure 8: The variation in NMW bite across counties, by occupation

(a) Professionals (b) Managers (c) Technical

(d) Admin (e) Skilled Trades (f) Services

(g) Sales (h) Process operatives (i) Elementary

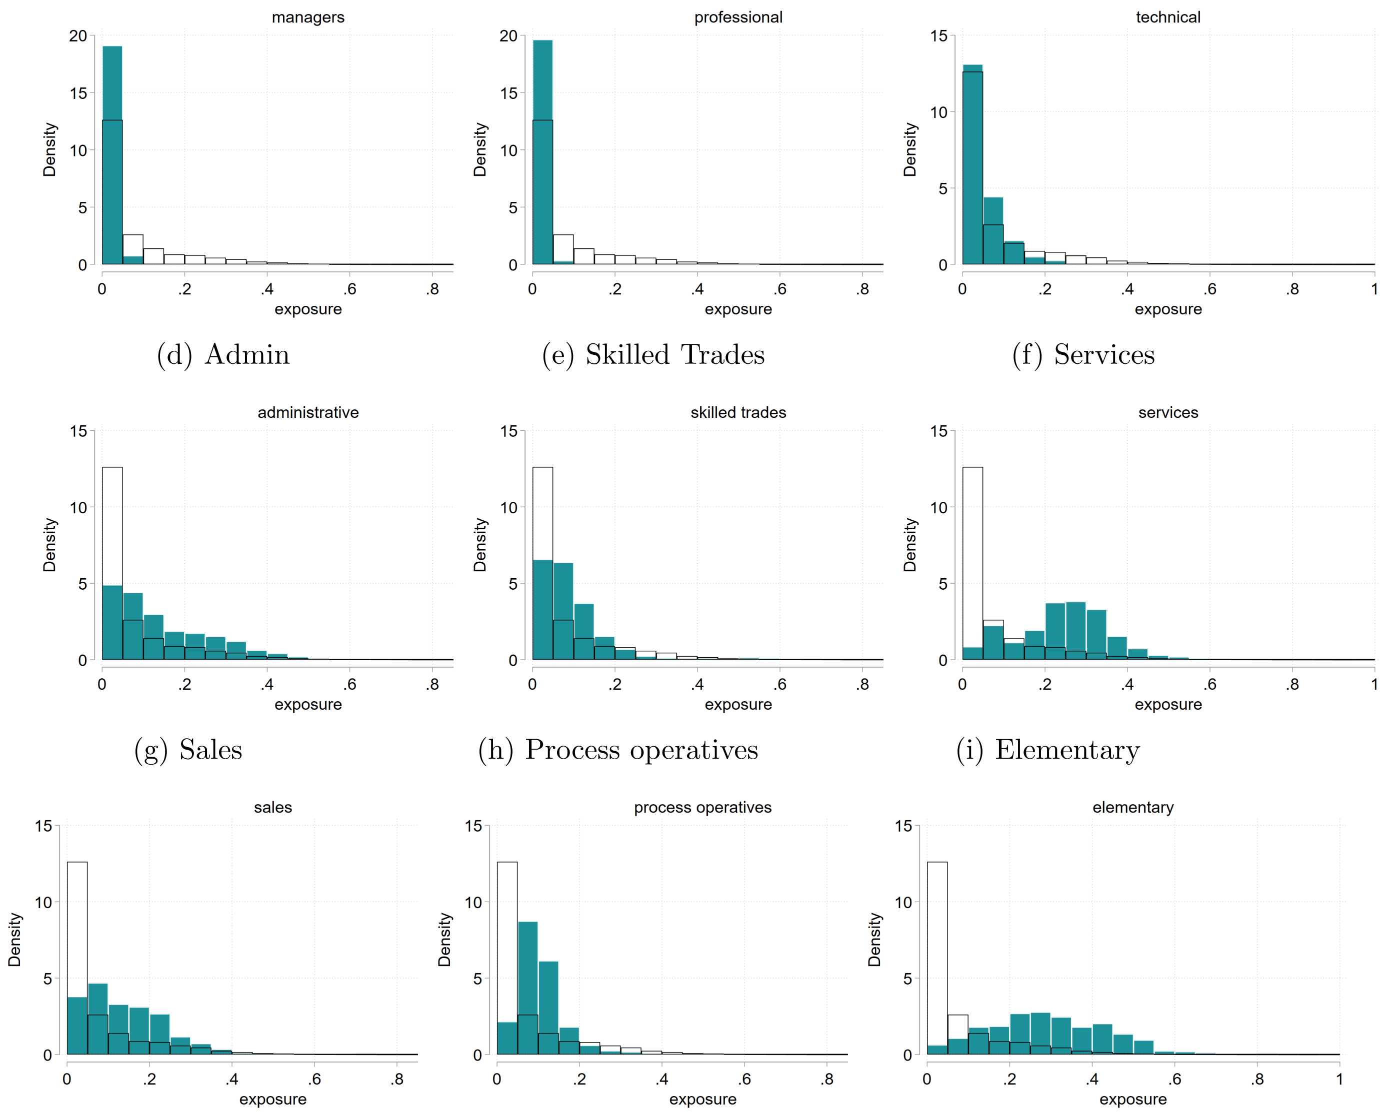

Notes: Each panel shows the distribution of the share of vacancies that are paid below the NMW (minimum wage bite) for all jobs (transparent bars) and or a given occupation (teal bars). Unit of observation is 3-digit SOC occupation by county.

4.3.1 Results

Our headline results on flexible work arrangements are summarised in Table 6. Panel (a) presents our headline estimate δ for flexible vacancies (column (1)), flexible and non-salaried vacancies (column (2)), and flexible and salaried vacancies (column (3)). The estimates show that the increase in the NMW increased the prevalence of flexible vacancies, but this increase was entirely driven by an increase in non-salaried flexible vacancies; as NMW rises, the share of salaried flexible jobs actually decreases. We confirm these diverging trends by re-estimating the regression for flexible work arrangements in the samples of non-salaried and salaried vacancies separately. The estimates in columns (4) and (5) show that while flexibility increased among non-salaried vacancies, the NMW increase had no impact on flexibility amongst the salaried jobs.

To further explore the mechanism, we re-estimate (2) for different levels of NMW bite. This confirms that occupation-county cells that were the most exposed to the minimum wage increase also saw the biggest adjustment in offered work arrangements. We split the occupation-county cells based on their NMW exposure: *low exposure* are cells with NMW bite in the third quartile of the NMW bite distribution (with average exposure os 12.7%), *medium exposure* corresponds to cells between the 75th and 90th percentile (with average exposure of 25.7%), and *high exposure* is defined as the top 10% of observations (with average exposure of 47.9%). The cells with NMW bite below the median are treated as the omitted group (with avergae exposure at 2.4%).

These results are summarised in panel (b) of Table 6. The estimated treatment effects increase in magnitude (and sometimes also statistical significance) as the exposure to NMW increases: the impact on the share of flexibility amongst non-salaried vacancies almost doubles as we move from the third quartile of NMW bite distribution to the top 10%. This difference is shown visually in Figure A.1 which plots the pre- and after-NMW prevalence of non-salaried flexible vacancies in administrative (low exposure) vs. elementary (high exposure) occupations. While there is little difference between the two groups before the introduction of NMW, non-salaried flexible vacancies become more and more prevalent in the high exposure group after the minimum wage increase, and this trend seems to accelerate over time.

Table 7 presents our difference-in-differences estimates for the remaining dimensions of work arrangements: duration of contract (permanent vs. temporary), number of working hours (full-time vs. part-time) and type of remuneration (salaried vs. non-salaried). There are two main takeaways. First, similarly to flexibility, rising labour costs have a significant impact on the prevalence of these work arrangements: jobs shifted away from "traditional" arrangements towards more temporary and non-salaried work. Second, the number of working hours demanded also dropped, but the size of this adjustment is an order of magnitude smaller that the adjustment in the demand for flexibility.

4.3.2 Spillover Effects

In the previous two sections, we showed that the increase in minimum wage resulted in a simultaneous increase in wages, and a shift away from traditional work arrangements. Next, we investigate which parts of the wage distribution were most affected by these changes in work arrangements: were the effects all concentrated at the NMW or did the change in work arrangements spills over to higher wage brackets otherwise seemingly unaffected by the minimum wage increase?

We adopt the following specification, adapted from Equation 4.6:

$$\frac{Y_{ocyqw}}{V_{ocyqw}} = \delta_w NMWbite_{oc} * After_{yq} + \alpha_c + \beta_{oyq} + \gamma_{ow} + \epsilon_{ocyqw} \tag{4.7}$$

where Y_{ocyqw} corresponds to the number of vacancies with a particular work arrangement in occupation o, county c, quarter q, year y and wage bin w, and V_{ocyqw} is the total number of vacancies in this *ocyqw* cell. The dependent variable hence captures the prevalence of work arrangement Y in the given wage bin. Other other variables are defined as in (2). We also add an occupation-wage-specific constant γ_{ow}. Standard errors are clustered at occupation-year-quarter cell, as in specification (2).

The raw estimates in Table A.2 show evidence of spillover effects beyond the wage bins[7] that were directly affected by the NMW introduction (as estimated in section 4.2). The increase in non-salaried flexible vacancies is not limited to the lowest two wage bins (£7-£8 and £8-£9): they also become more prevalent in the £9 - £10 bins and to a smaller extent in the £10 - £12 bin. The estimated treatment effects show these spillover effects can be positive as well as negative, i.e. go with or against the overall treatment effect. On one hand, the drop in permanent vacancies is strongest in the lowest wage bins, but permanent contracts become less prevalent in higher wage brackets too. On the other hand, part-time vacancies increase overall and in the lowest wage bins,

[7] We do not report the estimates for the lowest two wage bins, because they become implausible under the new minimum wage legislation. The denominator in these wage bins approaches 0, which would given rise to imprecisely measured dependent variable.

Table 6: Minimum Wage Exposure & Flexible Vacancy Share

	Flexible	All Flexible & Non-Salaried	Flexible & Salaried	Non-Salaried	Salaried
	(1)	(2)	(3)	(4)	(5)
Panel (a): Linear Specification					
Exposure x Treatment	0.1194***	0.1970***	-0.0695***	0.1541***	0.0206
	(0.0135)	(0.0113)	(0.0076)	(0.0127)	(0.0126)
R^2	0.2152	0.2211	0.0377	0.1705	0.0992
County F.E.	yes	yes	yes	yes	yes
SOC3 x Time F.E.	yes	yes	yes	yes	yes
Panel (b): Categorical Specification					
Control	0.0000	0.0000	0.0000	0.0000	0.0000
	(.)	(.)	(.)	(.)	(.)
Low Exposure	0.0139***	0.0161***	-0.0030**	0.0324***	0.0011
	(0.0025)	(0.0019)	(0.0014)	(0.0026)	(0.0017)
Medium Exposure	0.0390***	0.0492***	-0.0092***	0.0435***	0.0142***
	(0.0041)	(0.0033)	(0.0025)	(0.0035)	(0.0033)
High Exposure	0.0447***	0.0705***	-0.0238***	0.0601***	0.0122***
	(0.0043)	(0.0036)	(0.0025)	(0.0047)	(0.0043)
R^2	0.2153	0.2211	0.0377	0.1705	0.0993
County F.E.	yes	yes	yes	yes	yes
SOC3 x Time F.E.	yes	yes	yes	yes	yes

Notes: OLS regressions. Standard errors in parenthesis Standard errors in parenthesis and clustered at occupation-quarter level. * $p<0.1$, ** $p<0.05$, *** $p<0.01$. Dependent variable in columns (1)-(3) are dummy variables for whether a vacancy is schedule-flexible, schedule-flexible & non-salaried, and schedule-flexible & salaried respectively. Dependent variable in columns (4) and (5) are whether a vacancy is schedule-flexible in the sample of non-salaried and salaried vacancies respectively. In panel (b): low exposure are cells with NMW bite in the third quartile of the NMW bite distribution; Medium exposure corresponds to cells between the 75th and 90th percentile; High exposure is defined as the top 10% of observations. NMW bite below the median is the omitted group.

but decreases in higher wage bins.

In Figure 9, we focus further on the shifts the *distribution* of different work arrangements across the wage distribution. We normalise the estimated coefficients by demeaning them, so that the

Table 7: Minimum Wage Exposure & Other Vacancy Characteristics

	Permanent		Salaried		Full Time	
	(1)	(2)	(3)	(4)	(5)	(6)
Exposure x Treatment	-0.1343***		-0.1863***		-0.0651***	
	(0.0117)		(0.0149)		(0.0149)	
Control		0.0000		0.0000		0.0000
		(.)		(.)		(.)
Low Exposure		-0.0063***		-0.0096***		-0.0010
		(0.0022)		(0.0026)		(0.0024)
Medium Exposure		-0.0320***		-0.0508***		-0.0166***
		(0.0033)		(0.0039)		(0.0039)
High Exposure		-0.0495***		-0.0716***		-0.0140***
		(0.0040)		(0.0051)		(0.0048)
R^2	0.0567	0.0567	0.2399	0.2400	0.0520	0.0520
County F.E.	yes	yes	yes	yes	yes	yes
SOC3 x Time F.E.	yes	yes	yes	yes	yes	yes

Notes: OLS regressions. Standard errors in parenthesis and clustered at occupation-quarter level. * p<0.1, ** p<0.05, *** p<0.01. Dependent variable in groups (1)-(3) are dummy variables for whether a vacancy is permanent, salaried and full-time respectively. In panel (b): low exposure are cells with NMW bite in the third quartile of the NMW bite distribution; Medium exposure corresponds to cells between the 75th and 90th percentile; High exposure is defined as the top 10% of observations. NMW bite below the median is the omitted group.

overall change in the prevalence of a given work arrangement is fixed. Each panel in Figure 9 thus plots the estimated treatment effect on a particular work arrangement prevalence in a given wage bin, δ_w, relative to the average treatment effect across all wage bins.

These normalised estimates show the increase in the minimum wage leads to a divergence in work arrangements between the two lowest wage bins (£7 - £9) and the upper part of the wage distribution. When a certain type of work arrangement becomes less likely (permanent, full-time and salaried), low-wage vacancies become relatively much less likely to offer such an arrangement; for work arrangements that become more prevalent overall, they become *even more* prevalent low-wage vacancies. It is plausible that this polarising effect of NMW increase may have contributed to the divergence in work arrangements in the UK labour market described in section 3.

Figure 9: Relative NMW impact on different work arrangements across wage bins

(a) Permanent

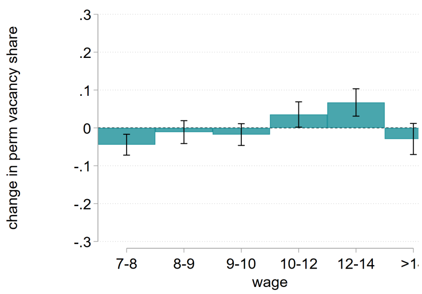

(b) Full-time

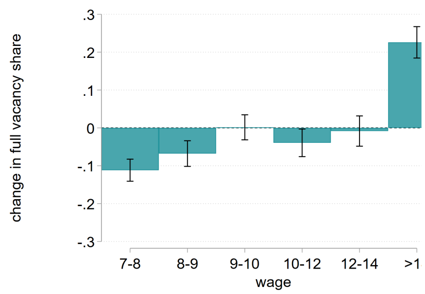

(c) Salaried

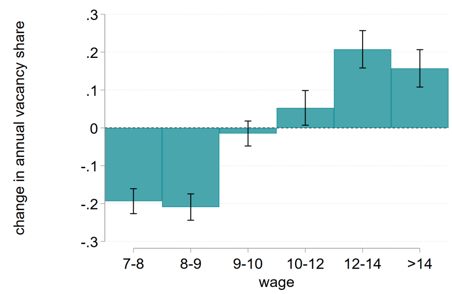

(d) Part-time

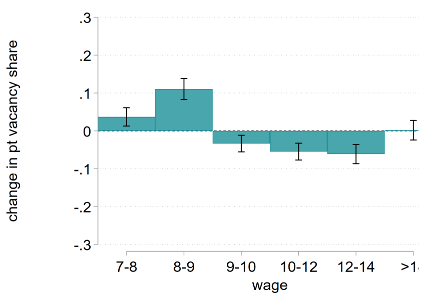

(e) Hourly pay

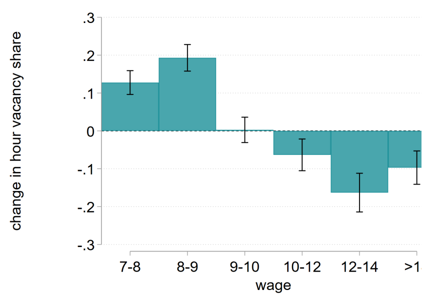

(f) Schedule-flexible

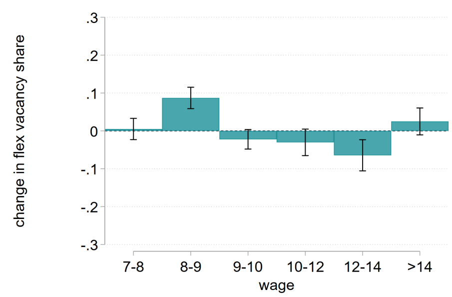

(g) Non-salaried & flexible

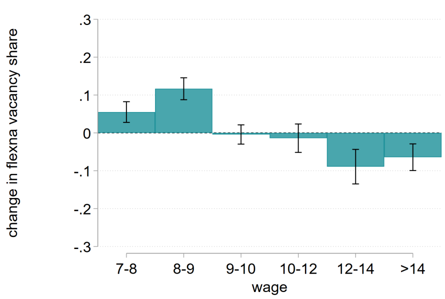

Notes: Normalised estimates of the treatment effect of NLW introduction on work arrangements in each wage bin. The estimates are taken from the specification given in Equation 4.7, normalised by the average treatment effect for the given work arrangement. Exposure to treatment is defined as % of vacancies in a given occupation-county unit which were paid below the new NLW before the intervention. Controlling for county, occupation-time and occupation-wage fixed effects.

5 Conclusion

In this paper we use a non-traditional dataset combined with a novel method to understand firm demand for flexible jobs in the UK. The wealth of our data – 46 million online vacancies posted between 2014 and 2019, classified by their work arrangements – allows us to analyse the variation in multiple dimensions of alternative work arrangements and along multiple lines. We document significant heterogeneity in employers' demand for non-traditional jobs across occupations, wage levels, and space; and a general upward trend in schedule-flexibility over the span of our dataset.

We argue that the implications of flexible jobs for worker welfare depends on the interaction between flexibility and other dimensions of the work arrangements. Our data shows that one should be careful to distinguish between salaried and non-salaried positions when analyzing flexible positions. In the latter, flexibility can give rise to earnings risk, while earnings are fixed with a salaried contract. We find that flexible and non-salaried vacancies are concentrated at low wages and are less likely to be permanent, in contrast with flexible and salaried positions. Non-salaried flexible vacancies grew at a faster rate than salaried ones, and their growth was concentrated in low-wage vacancies and occupations where the share of flexible jobs had already been relatively high.

We exploit the large and unexpected increase in the adult national minimum wage in 2016 to study whether firm demand for flexibility depends on labour costs. We estimate several empirical models to understand firm's reaction: we look at the shift in the distribution of posted wages, the overall change in the advertised work arrangements, and the distribution of these changes across wage brackets. We find that the higher national minimum wage led to an increase in the share of non-salaried flexible vacancies. This increase was predominantly clustered at low wage jobs just above the new minimum wage, leading to a polarisation in the work arrangements offered across the wage distribution.

Our paper has several implications for policymakers. In contrast to the existing literature (Mas and Pallais 2020), we show that alternative work arrangements advertised in the UK tend to cluster in low wage, non-salaried jobs, and that their share is rising over time. This presents a different picture to the narrative that flexible jobs are mostly prevalent among well-paid, highly-educated workers. In addition, our the response to a minimum wage increase suggests that policymakers need to be conscious of responses in contractual arrangements when designing policies aimed at supporting low-income individuals.

The paper also makes a contribution to the growing field of text analysis and machine learning in labour economics. We show that vacancy text can be used as source of information about various features of the offered job, our machine learning algorithm achieving a very good level of accuracy, precision and recall compared to simpler keyword searches. In practice, this approach allows us to study dimensions of the employer-employee relationship that are not captured well by traditional survey data, and thus somewhat unexplored in the literature.

References

Abraham, K. G. and A. Amaya (2019). Probing for informal work activity. *Journal of Official Statistics 35*(3), 487–508.

Adams, A., J. Prassl, et al. (2018). Zero-hours work in the united kingdom. *Conditions of Work and Employment Series* (101).

Adams-Prassl, A., T. Boneva, M. Golin, and C. Rauh (2020). Inequality in the impact of the coronavirus shock: Evidence from real time surveys.

Boeri, T., G. Giupponi, A. B. Krueger, and S. Machin (2020). Solo self-employment and alternative work arrangements: a cross-country perspective on the changing composition of jobs. *Journal of Economic Perspectives 34*(1), 170–95.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics 134*(3), 1405–1454.

Clemens, J., L. B. Kahn, and J. Meer (2018). The minimum wage, fringe benefits, and worker welfare. Technical report, National Bureau of Economic Research.

Clemens, J., L. B. Kahn, and J. Meer (2020). Dropouts need not apply? the minimum wage and skill upgrading. Technical report, National Bureau of Economic Research.

Datta, N. (2019). Willing to pay for security: a discrete choice experiment to analyse labour supply preferences.

Datta, N., G. Giupponi, and S. Machin (2019). Zero-hours contracts and labour market policy. *Economic Policy 34*(99), 369–427.

Deming, D. and L. B. Kahn (2018). Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics 36*(S1), S337–S369.

Duchini, E., S. Simion, and A. Turrell (2020). Pay transparency and cracks in the glass ceiling. *arXiv preprint arXiv:2006.16099*.

Forsythe, E., L. B. Kahn, F. Lange, and D. Wiczer (2020). Labor demand in the time of covid-19: Evidence from vacancy postings and ui claims. *Journal of Public Economics*, 104238.

Giupponi, G. and S. J. Machin (2018). Changing the structure of minimum wages: firm adjustment and wage spillovers.

Hall, J. V. and A. B. Krueger (2018). An analysis of the labor market for uber's driver-partners in the united states. *ILR Review 71*(3), 705–732.

Hershbein, B. and L. B. Kahn (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review 108*(7), 1737–72.

Javorcik, B., K. Stapleton, B. Kett, and L. O'Kane (2020). Unravelling deep integration: Local labour market effects of the brexit vote.

Marinescu, I. (2017). The general equilibrium impacts of unemployment insurance: Evidence from a large online job board. *Journal of Public Economics 150*, 14–29.

Mas, A. and A. Pallais (2017). Valuing alternative work arrangements. *American Economic Review 107*(12), 3722–59.

Mas, A. and A. Pallais (2020). Alternative work arrangements. Technical report, National Bureau of Economic Research.

Prassl, J. (2018). *Humans as a service: The promise and perils of work in the gig economy.* Oxford University Press.

Schütze, H., C. D. Manning, and P. Raghavan (2008). *Introduction to information retrieval*, Volume 39. Cambridge University Press Cambridge.

Taylor, M., G. Marsh, D. Nicol, and P. Broadbent (2017). Good work: The taylor review of modern working practices.

Wiswall, M. and B. Zafar (2018). Preference for the workplace, investment in human capital, and gender. *The Quarterly Journal of Economics 133*(1), 457–507.

A Additional Results

Table A.1: List of the positive and negative keywords and their weights based on the logistic regression model for our employment arrangements

Work arrangement	**positive keyword**		**negative keyword**	
Schedule flexible				
	flexible	2.0064	computer	-0.7026
	rota	1.6017	overall	-0.4603
	overtime	1.5349	suitable	-0.4288
Permanent				
	permanent	6.4692	temp	-2.8973
	typepermanent	6.2618	temporary	-2.6088
	type permanent	2.1474	fixed term	-1.5531
Full-time				
	full time	4.4086	part time	-3.7370
	hours	1.2398	class ntkn	-1.0573
	permanent full	1.1224	newly	-0.6960
Salaried				
	annum	2.8643	hour	-1.6584
	salary ntkn	1.6481	per hour	-1.3088
	year	1.1248	all jobs	-0.8721

Table A.2: The impact of NLW on share of vacancies of the given work arrangement within occupation-county unit, by wage

	(1) Permanent	(2) Flex	(3) Full	(4) PT	(5) Annual	(6) Hour	(7) Insecure
7-8	-0.1246***	0.0523***	-0.1816***	0.0080	-0.1989***	0.1758***	0.1075***
	(0.0140)	(0.0143)	(0.0149)	(0.0123)	(0.0168)	(0.0160)	(0.0140)
8-9	-0.0913***	0.1343***	-0.1377***	0.0816***	-0.2145***	0.2411***	0.1691***
	(0.0154)	(0.0144)	(0.0173)	(0.0142)	(0.0177)	(0.0179)	(0.0148)
9-10	-0.0978***	0.0250*	-0.0684***	-0.0625***	-0.0202	0.0508***	0.0482***
	(0.0147)	(0.0131)	(0.0168)	(0.0112)	(0.0168)	(0.0171)	(0.0130)
10-12	-0.0448***	0.0169	-0.1093***	-0.0837***	0.0475**	-0.0152	0.0384**
	(0.0170)	(0.0179)	(0.0186)	(0.0113)	(0.0234)	(0.0214)	(0.0191)
12-14	-0.0130	-0.0172	-0.0782***	-0.0902***	0.2022***	-0.1149***	-0.0367
	(0.0184)	(0.0210)	(0.0204)	(0.0129)	(0.0252)	(0.0260)	(0.0233)
>14	-0.1095***	0.0723***	0.1560***	-0.0271**	0.1519***	-0.0489**	-0.0118
	(0.0209)	(0.0181)	(0.0212)	(0.0132)	(0.0252)	(0.0225)	(0.0179)

Standard errors in parentheses

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

Notes: OLS regressions. Share of vacancies of a particular contract type in a given wage bin. Unit of analysis: 3-digit occupation-county-quarter-wage bin. Exposure to treatment is defined as % of vacancies in a given occupation-county unit which were paid below the new NLW before the intervention. Controlling for occupation-quarter FE, occupation-wage FE and county FE. Standard errors clustered at occupation-quarter level.

Figure A.1: The impact of NLW on share of income-insecure vacancies, occupations with different exposure

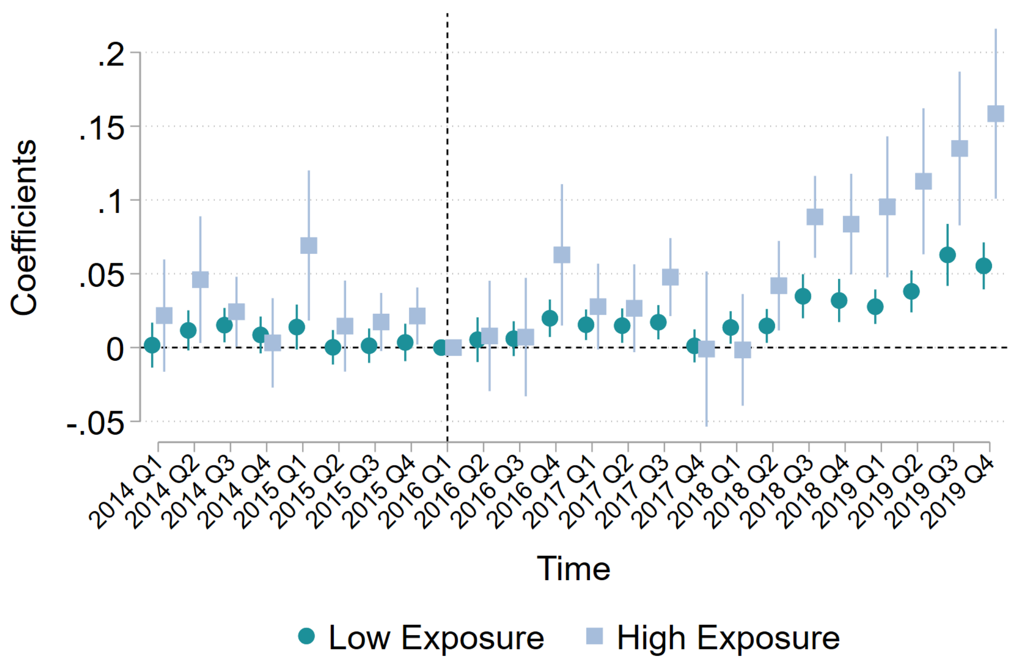

Estimates from an event study regressions on the share of vacancies that are income-insecure, before and after the introduction of NLW. Unit of analysis: 3-digit occupation-county-quarter. Exposure to treatment is defined as % of vacancies in a given occupation-county unit which were paid below the new NLW before the intervention. Controlling for occupation-quarter FE and county FE. Standard errors clustered at occupation-quarter level. Low exposure = administrative occupations. High exposure = elementary occupations. Data: Burning Glass vacancies for the UK, 2012-2019.