CESifo WORKING PAPERS

6766
2017

November 2017

Real Wages and Hours in the Great Recession: Evidence from Firms and their Entry-Level Jobs

Daniel Schaefer, Carl Singleton

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl
www.cesifo-group.org/wp

An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: www.CESifo-group.org/wp

CESifo *Working Paper No. 6766*
Category 4: Labour Markets

Real Wages and Hours in the Great Recession: Evidence from Firms and their Entry-Level Jobs

Abstract

Using employer-employee panel data, we provide novel facts on how real wages and working hours within jobs responded to the UK's Great Recession. In contrast to previous studies, our data enables us to address the cyclical composition of jobs. We show that firms were able to respond to the Great Recession with substantial real wage cuts and by recruiting more part-time workers. A one percentage point increase in the unemployment rate led to an average decline in real hourly wages of 2.8 per cent for new hires and 2.6 per cent for job stayers. Hours of new hires in entry-level jobs were also substantially procyclical, while job-stayer hours were nearly constant. Our findings suggest that models assuming rigid labour costs of new hires are not helpful for understanding the behaviour of unemployment over the business cycle.

JEL-Codes: E240, E320, J310.

Keywords: wage rigidity, Great Recession, hours worked, job-level analysis.

Daniel Schaefer
University of Edinburgh
School of Economics
Edinburgh / United Kingdom
daniel.schaefer@ed.ac.uk

Carl Singleton
University of Edinburgh
School of Economics
Edinburgh / United Kingdom
carl.singleton@ed.ac.uk

8 November 2017
Link to latest version: https://sites.google.com/site/danielschaefereconomics/research
We are grateful for helpful comments and advice from Steven Dieterle, Mike Elsby, Jan Grobovšek, Andrei Potlogea, Andy Snell, Andreas Steinhauer, Ina Taneva, Jonathan Thomas, Ludo Visschers, and participants at the 2017 Aarhus conference on Markets with Search Frictions. This paper is mostly based on the Annual Survey of Hours and Earnings Dataset (Crown copyright 2017), having been funded, collected and deposited by the Office for National Statistics (ONS) under secure access conditions with the UK Data Service (SN:6689). Neither the ONS nor the Data Service bear any responsibility for the analysis and discussion of the results in this paper.

1 Introduction

The Great Recession was the most severe economic contraction in the UK since the Second World War, yet employment declined less than in previous recessions. This resilience of employment has been attributed to flexible labour costs because aggregate real wages and working hours fell during the recent downturn (Crawford et al., 2013; Blundell et al., 2014; Gregg et al., 2014). But economy-wide averages tell nothing about the response of wages and hours *within* jobs, which is what determines firms' employment decisions in frictional labour markets. For example, suppose that wages within jobs are completely rigid, and workers switch from high- to low-paying jobs during recessions. In this case aggregate real wages would decline, even if firms' payments to employees are unchanged. In contrast to previous studies for the UK, we use a linked employer-employee dataset, which allows us to measure the response of real hourly wages and weekly hours worked *within* particular jobs.

Our main contribution is to combine the robust job-level measurement of responses in *real* wages and hours worked within the same methodological framework, for both new hires and job stayers. We present two novel findings: first, firms significantly reduced the real wages of new hires and job stayers within jobs during the downturn; second, the same firms kept the hours worked of job stayers unchanged, but significantly reduced the hours of new hires. A one percentage point increase in the unemployment rate leads to an average decline in real hourly wages of 2.8 per cent for new hires and 2.6 per cent for job stayers. Weekly hours worked of new hires decline by 1.5 per cent, but for job stayers remained nearly constant. A shift from full- to part-time work explains over half of the decline in the hours of new hires, however we see no significant difference between the wage responses of full- and part-time workers. This substantial flexibility for new hires could explain the relatively high job-finding rate during the UK's Great Recession.

In a wide class of labour market models, firms' employment decisions are forward-looking and dependent on expected labour costs. Therefore, we track new hires over three years in continuing matches to understand how persistent their initial real hiring wages and hours are. We find strong cohort effects: accounting for unobserved match quality, real wage growth came to a complete halt for cohorts hired during the Great Recession. For these employees, stagnant wages were only partially compensated for by larger increases in working hours with tenure on the job. These findings suggest that the sum of real wage payments in a job-match over time, i.e. the present value of labour costs for new hires, is more responsive to business cycle conditions than initial hiring conditions.

We use a simple empirical approach to measure the responses of real wages and hours to the business cycle. To obtain job-level measures of wages and hours, we first compute the median real wages and hours of new hires and job stayers in each job and year. We then estimate the semi-elasticity of these job-level measures to the unemployment rate, controlling for the changing composition of jobs over the cycle by including job-fixed effects. We identify a particular sample of jobs into which firms consistently hired before and during the recession. This matters, because we use within-job variation to measure responses to the Great Recession. If jobs with relatively rigid wages and hours simply stopped hiring during the recent downturn, then our sample would over-represent jobs with particularly flexible hiring conditions.

Our sample mostly consists of jobs with high turnover and low wages, which we call "entry-level". Since the employment of low-wage workers typically declines sharply during UK recessions (Blundell et al., 2014), it matters from an aggregate perspective why this group's employment did not drop more during the recent downturn. On average our entry-level jobs account for two-thirds of hires within their respective firms and a quarter of all new hires in the UK. These proportions increased during the Great Recession, both within firms and the whole economy, underlining the importance of these jobs in understanding the performance of the UK's labour market.

Since Solon et al. (1994) it has been recognised that the measured business cycle response of the average real wage typically underestimates the true response because of composition bias: the share in total hours worked of low-wage workers decreases during recessions, inducing a countercyclical bias. Several studies have used longitudinal data to address this bias (see the survey by Abraham and Haltiwanger, 1995). Most recently, Elsby et al. (2016) find that real wages are procyclical for UK employees working in the same job for at least one year, with an especially large response to the Great Recession. Pissarides (2009) however argues that what matters for a firm's hiring decision, and thus vacancy creation, are wages in new worker-firm matches. There exists some evidence that the wages of workers who change employers respond to the business cycle (see **?**Shin, 1994; Devereux, 2001, and Gertler et al., 2016 for US evidence). Devereux and Hart (2006) and Hart and Roberts (2011) find that the cyclicality of wages for British job changers significantly exceeds that of stayers.

But Gertler and Trigari (2009) explain why this worker-level evidence does not rule out wage rigidity within jobs: for example, if workers switch from high- to low-paying jobs during recessions, and vice versa during booms, then worker-level regressions will find a procyclical response of real hiring wages to business cycle conditions, even if wages within jobs are unchanged. But for firms' vacancy creation, and thus job-finding rates, this form of worker-level wage flexibility is not relevant. In frictional labour markets, firms will create vacancies so long as they expect a profit from doing so. This profit depends on expected revenue and the job-level hiring wages. Whether a new hire was previously employed at a higher or lower wage does not affect these costs.

If firms hire into jobs with relatively rigid wages in recessions, then the weighting in worker-level regressions is endogenous, and hence estimates are biased. Martins et al. (2012) propose measuring the cyclicality within jobs, by using the "typical" real wages of new hires in a case study of certain jobs. This approach trades off representativeness of the whole economy for confidence that economically meaningful responses can be estimated, at least from the perspective of what matters to firms. Martins et al. find that in Portugal real wages decrease significantly by 1.8 per cent when the unemployment rate increases by one percentage point. This is lower than our similarly obtained UK estimate of 2.8 per cent.

Our approach differs from Martins et al. (2012) in a subtle but important respect: in contrast to their study, we measure the responses of real wages and hours of job stayers using the same approach as for new hires. This provides a comparable benchmark value of real wage and hours flexibility, and allows us to assess whether these variables are especially flexible for new hires. Because we restrict attention to job stayers among the same firms as new hires, we can exclude firm-level differences as a source of bias when comparing the measured responses across the two groups of workers and jobs.

The institutional framework of labour markets is likely to affect the flexibility of wages and hours worked. Therefore, we further expand on Martins et al. by taking the effects of the National Minimum Wage on our estimates seriously. Intuitively, some entry-level jobs with relatively low hiring wages will be constrained in how they can adjust their wages downwards in response to the Great Recession. We compute counterfactual hiring wages using the method of DiNardo et al. (1996), and our findings suggest that the real wages of new hires would have declined a further 10 per cent if the National Minimum Wage had remained at its lower pre-crisis level.

Apart from the studies of Carneiro et al. (2012) and Martins et al. (2012) for Portugal, the only other estimates of hiring wage cyclicality at the job level are the findings of Stüber (2017) for Germany. He first measures the business cycle response of wages at the worker-level, controlling for unobserved worker and job heterogeneity. Stüber finds a semi-elasticity of average daily real wages to the unemployment rate of 1.3, which does not significantly differ between new hires and job stayers. Because of the potential problems of endogenous sample selection when running worker-level regressions, he also estimates a job-level version. His findings suggest that job-level real daily wages of new hires and job stayers decline by 0.9 per cent if the unemployment rate increases by one percentage point in Germany. But, unlike the approach of Martins et al. (2012), he does not fix a particular sample of jobs. Again, if firms stop hiring into jobs with relatively flexible wages, and thus only rigid job-level real hiring wages are observed, then his results will be biased towards finding smaller semi-elasticities.

As Stüber explains, the method applied by Martins et al. and us is not immune from composition bias: within jobs it is likely that relatively low-skilled and low-wage workers are the first to become unemployed in recessions. This induces a countercyclical bias, and so our results provide lower bounds for the response of hiring wages within jobs. However, because the jobs in our entry-level sample are mostly low-skilled and concentrated in the hospitality and trade services industries, we expect that changes in the composition of workers within jobs do not substantially affect typical job-level wages. This is in line with Yagan (2017), who argues that, conditional on firm-fixed effects, the composition of workers and the tasks performed by them does not vary notably in the US retail industry.

All of the aforementioned studies focus on real wages. But firms can also adjust their labour costs by decreasing the hours worked per employee. Our detailed dataset allows us to also examine the responses of weekly hours worked for the same employees and jobs, which previous studies could not address. We present a novel and robust finding that firms responded to the Great Recession by significantly decreasing hiring hours. We find that a shift from full- to part-time explains half of the overall decline in hiring hours within jobs. However, hours worked were not responsive to the Great Recession for job stayers.

2 Measuring how real wages and hours responded to the Great Recession

To measure the response of real hourly wages and weekly hours to the Great Recession we use a two-step regression approach (Solon et al., 1994; Martins et al., 2012). Compared to the alternative one-step approach, the results are more transparent and we do not have to rely on asymptotic theory to obtain robust estimates of standard errors. We expand on our choice of method further below. In the first step, for hiring wages we use least squares to estimate

$$w_{jt} = \alpha_j + \beta_t + \boldsymbol{x}'_{jt}\boldsymbol{\delta} + \varepsilon_{jt} \, , \tag{1}$$

where w_{jt} is the median log real hiring wage in the 4-digit occupation-firm pair j (hereafter job j) and period t. We include job-fixed effects α_j and period-fixed effects β_t. The error term ε_{jt} gives the remaining heterogeneity in w_{jt} which is not job- or period-specific, after controlling for time-varying job characteristics in the vector $\boldsymbol{x}_{jt}$. The baseline set of covariates for new hires at the job level are: a cubic function in age and firm size, the share of female employees and the share of employees covered by a collective agreement. Our results remain virtually unchanged when we instead use dummies for ranges of these variables. We include these covariates to control to some extent for changes in the composition of employees within jobs over the business cycle.

The parameter estimates $\hat{\beta}_t$ from (1) are a series of period-means of log wages, regression adjusted for changes in the composition of jobs in the sample. In the second step, we relate this series to the Great Recession by regressing it on the unemployment rate U_t:

$$\hat{\beta}_t = c_0 + c_1 t + \gamma U_t + e_t \; . \tag{2}$$

We vary the specification of both steps for robustness, but the baseline second step includes a constant and a linear time trend. We measure the response of real wages to the Great Recession by the coefficient estimate $\hat{\gamma}$, the semi-elasticity of real wages with respect to the unemployment rate. If instead we regressed job-level wages directly on the unemployment rate, then errors would be cross-sectionally correlated. This is because the cyclical indicator does not vary across jobs: usual standard errors would underestimate the uncertainty of coefficient estimates.[1] Therefore, following the recommendations of Donald and Lang (2007) and Angrist and Pischke (2009), we use a two-step procedure on within-period averages. This approach is transparent and standard error estimates are more reliable than estimating a covariance matrix that is robust to cross-sectionally correlated errors (or "cluster robust") with relatively few periods.

For job-stayer wages we alter the first-step regression. Let w_{qk} be the median wage of job stayers in some job q, which is specific not only to some occupation-firm pair, as per j above, but is also specific to two consecutive periods: i.e. job stayers observed between years 1998-9 and 2008-9 who work in the same occupation-firm pair j would have different values for q. Whether the wage refers to stayers in the first or second consecutive period is indicated by k, equal to zero or one respectively. We thus account

[1] For an illustration, note that the error term v_{jt} of the one-step regression

$$w_{jt} = \alpha_j + c_1 t + \gamma U_t + \boldsymbol{x}'_{jt}\boldsymbol{\delta} + v_{jt}$$

consists of a job-specific component ε_{jt} and a period-specific term e_t, such that $v_{jt} = \varepsilon_{jt} + e_t$. The error term v_{jt} is cross-sectionally correlated because of e_t, which is common across all jobs j within a period.

for wages rising with tenure, using least squares to estimate

$$w_{qk} = \alpha_q + \beta_{T(qk)} + \lambda k + \boldsymbol{x}'_{qk}\boldsymbol{\delta} + \varepsilon_{qk} \ , \tag{3}$$

where $T(qk)$ is a function indicating that a job q is observed in period t, and $\boldsymbol{x}'_{qk}$ contains time-varying characteristics: age squared and cubed, tenure in the firm squared and cubed, firm size and its square, and the share of employees covered by a collective agreement. Linear terms for age and tenure are omitted as these would be collinear with k, and the average effect of these variables is controlled for by the estimated linear trend $\hat{\lambda}$. Although rewriting and estimating (3) in first differences over k is a more intuitive representation of how we estimate job-stayer wage cyclicality, we proceed with the equation in levels to obtain more directly comparable estimates of $\hat{\beta}_t$. The second step regressions for job-stayer and hiring wages are identical.

To measure the response of hours worked we estimate the same two-step models, replacing the dependent variables in (1) and (3).

2.1 The Annual Survey of Hours and Earnings and other data used

The Annual Survey of Hours and Earnings (ASHE), 1997-2016, is based on a one per cent random sample of employees, drawn from HM Revenue and Customs Pay As You Earn (PAYE) records. A small number of workers not registered for PAYE, who tend to receive very low pay, either due to low hourly wages or hours worked, or both, are not sampled. Questionnaires are sent to employers, who are legally required to complete them with reference to payroll for a period in April. The ASHE is generally considered to provide accurate records of pay components (Nickell and Quintini, 2003).

The dataset can be viewed as a panel of employees without attrition, forming an approximate one per cent random and representative sample of UK employees in every year.[2] Particularly valuable for our analysis are the longitudinal identifiers for

[2]The two main reasons why an individual might not be observed in some year are: either being truly non-employed, or having changed employer between January and April. Since the survey questionnaires

individuals (1997-2016) and enterprises (2003-2016). We use the terms "firm" and "enterprise" synonymously. The latter in this case is a specific administrative definition of UK employers, which could contain several local units (or plants). We believe this is the appropriate level to study firm- or job-level wages, because in most organisations pay-setting practices are determined at the enterprise level.

In another paper we use a combination of the exact number of employees, 4-digit industry classification and information on the legal status of an enterprise to define the boundaries of larger firms within each year back to 1996 (Schaefer and Singleton, 2017). But, here we want to link firms longitudinally before and after 2003. Therefore we impute values of the enterprise identifier backwards from 2003, using consecutive observations of individuals who have not changed jobs between years, as well as using employment start dates. We then use the available within-year employer local unit identifiers to impute more enterprise values. For further information on how we construct this employer-employee panel from ASHE cross-sections, and other adjustments made to the data and the sample selection, see Appendix A.

Our analysis focuses on two main variables: basic weekly paid hours and the hourly wage rate, which equals the ratio of gross weekly earnings to the former, all excluding overtime. We refer to these simply as hours and wages. Monetary values are deflated using the Consumer Price Index (CPI).[3] For comparability with other studies, we include some statistics about nominal wage changes in Appendix E. We consider working-age employees (aged 16-64) in the private sector, who have non-missing records of earnings and hours. We include only the main job observation of an individual, which must not be at trainee or apprentice level, and not have incurred a loss of pay in the reference period for whatever reason. To avoid some spurious hourly wage rates we only keep observations with 1-100 basic paid weekly hours. Since the data are not top-coded, we drop the highest one per cent of weekly or hourly earners.

are in most cases sent in April to the employer's registered address from January PAYE records, workers who switch employers during these months are undersampled.

[3]For robustness we also compute results using the Retail Price Index (RPI). All prices were obtained from UK National Statistics, accessed 24/4/2017.

Our main indicator variable for the Great Recession is the working-age unemployment rate: the number of people unemployed divided by the economically active population.[4] To correspond with the timing of the ASHE, we use average values over the previous four quarters for all price series and business cycle indicators. For example, an estimate of the 2009 wage for new hires is compared with the average unemployment rate over the preceding twelve months, when those hires would have been made. We use the unemployment rate for comparability with the wider literature. In Section 4 we discuss the robustness of our main results to this choice.

2.2 Constructing the baseline sample of entry-level jobs and their firms

We create a sample of entry-level jobs following Martins et al. (2012), applying similar selection criteria. We first restrict our sample to observations for the years 2003-16, because for this period we have almost complete records of firm identifiers and employment start dates. We exclude all firms which are observed for less than three years. Jobs are defined at the 4-digit occupational level within firms (for example, "Housekeeper" vs. "Waiter or waitress" in a hotel), whereby the same occupation in two different firms is treated as two separate jobs. We define a new hire as any employee with less than one year of tenure with a firm.

For a job to be defined as entry-level, we require at least three observations of new hires in a year, and this must be the case for the job in at least half of the years when the firm is observed in 2003-16. Recalling that the ASHE is an approximate one per cent random employee sample, these requirements impose an effective lower bound on firm size in our entry-level jobs samples. Of the firms in our baseline sample, 95 per cent have more than five hundred employees. After identifying entry-level jobs over 2003-16, we add further observations of new hires in these jobs back to 1998. These earlier hires in the sample tend to be older individuals and subsequently have longer tenure with the firm, a result of how we recursively impute firm identifiers before 2003.

[4]Source: ONS Labour Market Statistics, April 2017, available at https://www.ons.gov.uk/.../apr2017; accessed 24/04/2017.

We do not claim that this sample represents all entry-level jobs in the economy, nor that the firms always hire into the same jobs. Instead the analysis of wages and hours in this sample should be viewed as a case study, where we do what is possible to control for composition bias in hiring over the economic cycle: we only study the real wages and hours of new hires in jobs where we can observe at least some hiring regardless of the economic cycle. In what follows we refer to the sample of firms which have these jobs as consistent-hiring-firms (CH-firms). The selection criteria are naturally somewhat arbitrary, though hopefully reasonable. We vary them for robustness when discussing our main empirical results.

2.3 Summary of new hires, job stayers and entry-level jobs

The entry-level jobs sample consists of 309 firms hiring into 391 jobs (Table 1). Our sample is unbalanced since some jobs are not observed in all years during 1998-2016. As Martins et al. (2012) note, the most important consideration is that the number of entry-level jobs should not vary systematically over the business cycle, as this would result in endogenous sample selection. The contemporaneous correlation of the number of entry-level jobs in our sample and the unemployment rate is insignificant (p-value: 0.49), and no other cyclical patterns are evident in Table 1 column (1). The median number of new hires per entry-level job is seven over the sample period.

A contribution of this paper is that we analyse the real wages and hours of job-stayers within firms which have at least one entry-level job. Job stayers are employees who are still working in the same occupation-firm as in the last reference period, hence we exclude the effects of cyclical job-switching into better or worse matches. We include only jobs with at least three job stayers in at least half of the years when the firm is observed during 2003-16. The sample consists of 7,779 repeated observations of occupation-firm pairs, totalling 158,194 job stayers. The selected jobs represent on average nearly 90 per cent of all job stayers in the CH-firms sample over the whole period.

New hires in the CH-firms sample are younger, more likely to be female, and less likely to work full-time than job stayers (columns (1) and (2), Table 2). The wages and

TABLE 1: Number of new hires, entry-level jobs, and consistent-hiring-firms by year

Year	New hires (1)	Entry-level jobs (2)	Firms (3)	Unemployment rate (4)
1998	948	116	93	6.80
1999	1,244	139	113	6.28
2000	1,358	148	113	5.94
2001	2,496	198	152	5.33
2002	2,821	219	180	5.15
2003	2,319	234	183	5.22
2004	2,460	252	191	4.98
2005	3,802	290	224	4.78
2006	3,502	289	225	5.02
2007	3,499	294	225	5.55
2008	3,609	289	221	5.34
2009	3,414	272	213	6.23
2010	2,781	258	203	7.96
2011	3,254	276	213	7.99
2012	3,178	262	206	8.37
2013	3,221	249	193	8.09
2014	3,374	262	206	7.48
2015	3,890	262	203	6.04
2016	3,507	242	186	5.39
Total	54,677	4,551	3,543	
Unique	48,744	391	309	

Notes.- age 16-64, private sector only. Source of the unemployment rate series is discussed in Section 2.

TABLE 2: Descriptive statistics for employees: comparison of the consistent-hiring-firms sample and the whole ASHE (all firms and jobs), 1998-2016

	CH-firms		ASHE	
	New hires (1)	Job stayers (2)	New hires (3)	Job stayers (4)
Mean age (years)	28	37	32	41
Female share	0.57	0.52	0.47	0.42
Full-time share	0.36	0.70	0.66	0.79
Median real hourly wage	5.24	7.04	6.29	8.43
Median basic weekly hours	21.6	36.0	36.5	37.4
Median real weekly earnings	117	260	225	313
Median firm size (n. of employees)	6,588	6,588	45	29
Firm size growth (p.a.)	4.3%	4.3%	7.9%	7.9%
N (000s)	55	158	222	1,307

Notes.- age 16-64, private sector only. Monetary values in GBP, deflated to 1998 prices using CPI. Descriptives for job stayers refer to their latter longitudinally linked observations.

basic hours of new hires are lower than for job stayers. The same statements hold for the entire ASHE, (columns (3) and (4), Table 2), though the difference between the hours worked by all new hires and job stayers is considerably smaller: almost two-thirds of hires into entry-level jobs are part-time. When compared with the whole economy, the lower average age, real wages, and basic hours in the CH-firms sample can be explained by differences in industry and occupation composition. Over two-thirds of new hires are made by firms in the "accommodation and restaurant" and the "industrial cleaning and labour recruitment" industry. Similarly, the largest shares are employed as service or sales workers (see Tables C1-C2 for the complete industries and occupation breakdowns). Firms in these industries account for approximately a third of all employees in the private sector (Jäger, 2016). This is also reflected in larger firms dominating the CH-firms sample. These large firms have average annual growth in the number of their employees of around four per cent, while the average for all firms is around eight per cent. Although the observable characteristics of new hires in entry-level jobs exhibit secular trends during 1998-2016, we do not see any notable cyclical patterns (Appendix Figure D1).[5]

Since our subsequent analysis takes place at the job level, we compute median wages and hours within jobs each year (our results do not change significantly when we use average wages and hours instead). Figure 1 shows the distribution of real wages of new hires after subtracting their respective entry-level job median wages.

Although we can expect some dispersion around these median wages, the robustness of our econometric approach and the meaningful interpretation of any results to some extent depends on us capturing "typical" hiring wages. Some dispersion around the median hiring wages, indicated by zero, is visible in Figure 1. More than 50 per cent of hiring wages lie within a range of five log points, and almost 90 per cent within 10 log points around the job-specific median. The dispersion around the typical hiring wage is approximately constant over time and does not vary systematically with the business

[5]Appendix Figure D1 shows that the share of men among new hires increases steadily by around ten per cent from 1998 to 2016, while the share of full-time employees decreases. As a consequence of our recursive sample construction prior to 2003, the average age of hires decreases by over five years from 1998 to 2003. Including controls for the average age of hires within a job in the following analysis does not change our results.

FIGURE 1: Distribution of differences between log real wages of new hires and their median values within entry-level jobs, 1998-2016

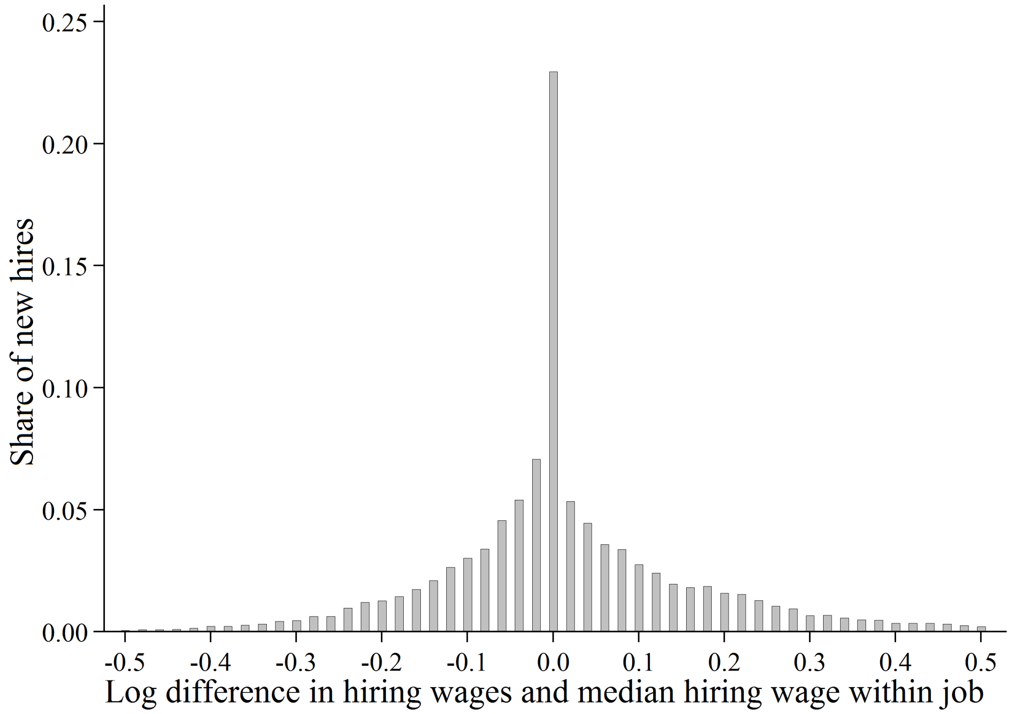

Notes.- within each entry-level job and year, the median hiring wage is subtracted from all hiring wages in that job and the resulting log differences are collected in bins with a width of two log points.

cycle. In particular, the mass in the tails of the distribution does not change and the inter quartile range is constant over the sample period. In the robustness discussion below we show that our results do not notably change if we use mean wages or hours within jobs instead.

3 Main results: Estimated job-level responses to the Great Recession

Figure 2 shows the estimated time series $\hat{\beta}_t$ for new hires and job stayers in CH-firms from regression models (1) and (3) over the period 1998-2016. The short-dashed lines indicate 95 per cent confidence intervals for the point estimates of $\hat{\beta}_t$ for new hires, using standard errors robust to serial correlation at the job-level. The confidence intervals for job stayers (not shown) are very narrow and lie within the intervals for new hires, except

FIGURE 2: Estimated period-fixed effects for log real wages and log hours worked, including 95% confidence intervals for new hires, 1998-2016.

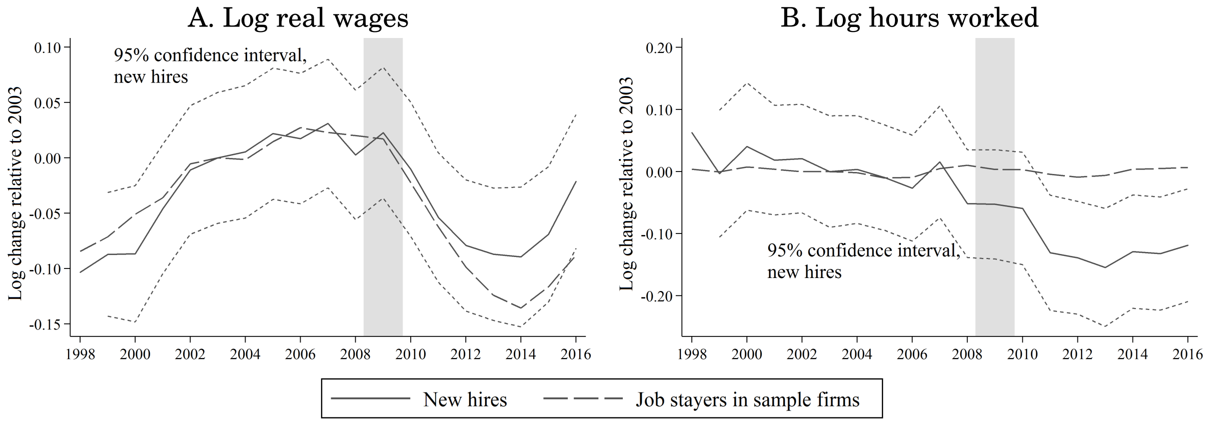

Notes.- the 95% confidence interval for job stayers (not shown) is very narrow and lies within the interval for new hires except in 2015-16 for real wages and 2011-16 for hours worked. Standard errors are robust to clustering at the firm-level. Excluded reference category in first-step regression (1) is 1998 for new hires, and regression (3) excludes 1998 & 2016. Series-specific linear trends removed from panel A. Series normalised to zero in 2003. Shaded area marks official UK recession dates. "New hires" are for wages in entry-level jobs where employees have less than twelve months of tenure. "Job stayers in sample firms" are for jobs and employees who have tenure greater than twelve months, and only for firms which are ever represented in the CH-firms sample.

in 2015-16 for real wages and 2011-16 for hours worked. All series are normalised to zero in 2003 and series-specific linear trends have been removed from the series for wages for comparability. We estimate (1) with the unbalanced baseline panel of jobs described in Table 1, using period and job dummies. Thus, these time series should be interpreted as composition-adjusted real wages and hours of new hires and job stayers. In Appendix Figure D5 we show the series for real wages without removing their respective linear trends.

Panel A shows that real hiring wages increased above trend by around 10 log points between 1998 and 2008, similar to job stayers in the same firms. During the Great Recession hiring wages remained below trend by around 10 log points until 2014, before slightly recovering over the next two years. The real wages of job stayers plummeted by almost 15 log points during the downturn, relative to trend. For comparison, Elsby et al. (2016) document a decline in job-stayer real wages in the whole economy between 2008 and 2012 of 14 log points for men and eight log points for women. Panel B shows the estimated series for hours worked among the same employees, jobs and time period. Hiring hours decreased by over 10 log points between 2007 and 2012, being approximately constant before and after. In contrast, the hours worked by job stayers saw no significant change during the Great Recession.[6]

We measure the response of real wages and hours to the Great Recession by estimating the second-step regression (2) using least squares. As recommended by Solon et al. (2015), we do not use weighted least squares (WLS) in our baseline regressions, because the least squares residuals do not display significant evidence of heteroskedasticity. Nevertheless, we later compare estimates from our baseline to those obtained using two different WLS estimators: (1) weights equal to the number of new hires; (2) weights equal to the number of entry-level jobs. Estimates from the first WLS estimator suffer from endogenous sample selection, because hiring volume is likely to depend on the cyclical response of wages and hours. However, the sign of the induced bias is informative about the bias in other studies using worker-level data. The second estimator is the WLS procedure applied by Martins et al. (2012), which accounts for the varying sample sizes of entry-level jobs over time.

The first row of Table 3 displays the main (or baseline) results, measuring the semi-elasticity with respect to a one percentage point (p.p.) increase in the unemployment rate: real hourly wages of new hires and job stayers decrease by 2.8 and 2.6 per cent if the unemployment rate increases by one p.p.. These estimates are significantly different from no response, but do not significantly differ from one another: Appendix Table B3 shows the coefficient estimates when we regress the difference between new hires and job stayers in the estimated series of $\hat{\beta}_t$ on the unemployment rate, using our two-step approach. In column (3) we find that hiring hours respond by around 1.5 per cent, compared to only 0.2 per cent for job stayers in column (4). The significant decline in average hours per worker during the Great Recession has been discussed before (Blundell et al., 2014; Pessoa and Van Reenen, 2014; Borowczyk-Martins and Lalé, 2017), though not at the job level. Blundell et al. (2008) find that UK workers adjust hours worked in response to welfare reforms usually by changing firms, and

[6] Appendix Tables C3-C4 display the underlying values of all series in Figure 2.

this is particularly true for larger firms and in the services industry. To the best of our knowledge, the relatively greater and large response of hiring hours has not been documented previously.

TABLE 3: Estimated semi-elasticity of real wages and hours with respect to the unemployment rate, 1998-2016

	Wages		Hours	
	New hires (1)	Job stayers (2)	New hires (3)	Job stayers (4)
1. Baseline	-2.83^{***}	-2.60^{**}	-1.47^{***}	-0.20
	(0.87)	(1.13)	(0.42)	(0.22)
2. Including controls for share of full-time workers	-2.78^{***}	-2.71^{**}	-0.68^{***}	-0.04
	(0.88)	(1.17)	(0.26)	(0.19)
3. Job hires in at least 25% of years when firm is observed	-2.44^{***}	-2.61^{**}	-0.43	-0.16
	(0.85)	(1.11)	(0.28)	(0.13)
4. All jobs observed in at least 2 years	-2.48^{***}	-2.90^{**}	-0.47^{**}	-0.16
	(0.86)	(1.16)	(0.18)	(0.10)
5. Baseline sample, but weighted by number of employees per year	-2.15^{***}	-1.88	-2.72^{***}	-0.43
	(0.64)	(1.03)	(0.68)	(0.22)

Notes.- second-step regression results, estimates $\hat{\gamma}$: responses of the period-fixed effects $\hat{\beta}_t$ to the unemployment rate; regression specifications as in (1)-(3). First row refers to the main/baseline estimates. Second row includes an additional time-varying control for the share of full-time workers in a job. Third row changes the selection criteria for entry-level jobs, such that they have to be fulfilled in at least a quarter of years when the firm is observed, instead of a half, and those firms have to be observed for at least five years. Fourth row includes all job observations which hire in at least two years, or with at least two consecutive years of observations for job stayers. Fifth row uses WLS in the first-step, with weights proportional to the number of new hires or stayers in each job.
Newey-West standard error estimates robust to first-order serial correlation in parentheses.
*** Statistically significant at the 1% level; ** at the 5% level, two-sided tests.

We also find that real weekly earnings (excl. overtime) of new hires decline by 4.7 per cent if the unemployment rate increases by one p.p., while job-stayer earnings decline by 2.9 per cent. Because the covariance between wages and hours is positive, these estimates exceed the sum of the corresponding values in the first row of Table 3. We re-estimate regression (1), but in addition control for changes in the share of full-time workers within jobs. The second row of Table 3 shows that for hiring hours the semi-elasticity estimate falls to 0.7: over half of the recessionary decrease in hiring hours can be attributed to a shift from full- to part-time hiring. However, the response of real wages to the unemployment rate does not differ significantly between full- and part-time hires and job stayers.

In the third row of Table 3 we include jobs which hire less frequently, increasing the sample number of entry-level hires by 25 per cent. We find a slightly smaller response of hiring wages to the recession, while the response of hours becomes insignificant. This suggests that not keeping the sample of jobs fixed induces a countercyclical bias in our estimates. When we create a balanced panel, by considering only jobs which fulfil our selection criteria in all years 2003-16 (not shown in Table 3), we find semi-elasticities of hiring wages and hours of 2.6 and 1.3 respectively. Therefore, our composition-adjusted baseline sample is relatively unaffected by selection bias over the business cycle.

As previously explained, the approach of Martins et al. (2012) has the advantage that we can be confident that what we measure is the response of real wages and hours of new hires in entry-level jobs. The potential disadvantage is that this response may not be representative of the whole economy. When we additionally include all other jobs in the ASHE for which we observe hiring in at least two periods (and similarly for job stayers), then the estimated time series of period-fixed effects from the first step resembles the series from our baseline sample (Appendix Figures D3 and Figures D4). The exception is hiring hours, which are less cyclical at the job level in the whole economy. This sample contains over twice as many hires and six times as many job stayers as our baseline, and the fourth row in Table 3 shows comparable estimates to our baseline. Real wages of all new hires and job stayers in this sample decrease by 2.5 and 2.9 per cent, respectively, for each p.p. increase in the unemployment rate. However the response of hiring hours is less pronounced in this larger but more representative sample, though still significant at the five per cent level. Thus, jobs with less flexible hiring wages and hours were more likely to stop hiring altogether during the Great Recession, inducing a countercyclical selection bias. In contrast, the selection bias for real wages of job stayers seems to be procyclical.

The final row of Table 3 shows values for $\hat{\gamma}$ when we estimate the first step using WLS, with weights proportional to the number of employees in each job. In this case the semi-elasticity estimate for hiring wages of 2.2 is the smallest of all the specifications described here. Therefore, weighting jobs by the number of hires induces a countercyclical bias in the estimated response of real wages to business cycle conditions: relatively more hires are made in jobs with relatively rigid hiring wages. Jobs which hired relatively more employees than others during the downturn also decreased the hours worked per hire more: the response to a one p.p. increase in the unemployment rate for entry-level hiring hours is more than 85 per cent larger than without endogenous weighting. Therefore, weighting by hiring volume overestimates the responsiveness of hiring hours. Not surprisingly, the firms who hire relatively more during recessions are also able to move their workforces towards shorter hours and part-time working. When Stüber (2017) similarly weights his job-level regression, German real hiring "wages" become more procyclical, seemingly contradicting our findings. But the results here could explain why we reach an opposite conclusion on the direction of bias induced by the endogenous hiring volume of jobs, which is also more in line with the hypotheses by Gertler and Trigari (2009) and Martins et al. (2012): the German data offers information on annual earnings and the number of days worked, and so Stüber's measures of real wages are better understood as average daily earnings. As our estimates show, the procyclical bias in hiring hours is large, and exceeds the countercyclical bias in hiring wages. When combined, this can cause a procyclical bias in earnings.

Finally, we make the sample selection criteria more exclusive, by increasing the required minimum number of hires in a job per year for it to be included as an entry-level job. The estimated semi-elasticity of hiring hours significantly increases in the number of minimum hires (see column (3) of Appendix Table B2), nearly doubling when we require at least 10 hires per job and year. The estimated semi-elasticity of real hiring wages slightly increases in absolute terms, peaking at 2.9 when we raise the minimum number of hires to 7 per job and year. Varying the minimum number of hires generally affects the measured response to the Great Recession of both wages and hours. However, our main finding is unchanged: real wages of new hires are marginally more responsive to the unemployment rate than wages for job stayers in the same firms. Similarly, hiring hours always respond more strongly to the Great Recession than job-stayer hours.

4 Robustness and further discussion

The main results described above show that UK firms were able to significantly decrease the real labour cost per employee in response to the Great Recession. To address robustness, in this section we apply alternative estimation procedures. We further discuss the measurement of real wage cyclicality, as well as the wider implications of these results. All of the additional analysis here uses the baseline consistent-hiring-firms sample of employees and jobs.

4.1 Using other specifications of the regression model

Table 4 displays results from varying the specification of the second-step regression (2), while the first step remains unchanged. The main baseline results are repeated in the first row. The second row shows that when we include a quadratic time trend, wages decline marginally less when the unemployment rate increases, but the hours responses are approximately unchanged. We prefer to only include a linear trend because of the small number of periods in our dataset.

We also re-estimate (2) in first differences, to address potentially spurious estimates if wages, hours, or the unemployment rate are integrated. As in our baseline results, the real wage growth of new hires and job stayers does not respond significantly differently to changes in the unemployment rate: if the change in the unemployment rate increases by one p.p., then the growth of hiring wages decreases by 1.6 per cent, and for job-stayers' wages by 1.8 per cent. This is comparable to the finding by Devereux and Hart (2006) of 1.7 per cent for all job stayers in the UK during 1975-2001.

For comparability with Martins et al. (2012), we re-estimate (2) using WLS, with weights proportional to the number of jobs per period in the first step. The resulting estimates in the final row of Table 4 are qualitatively unchanged from the baseline. However real wages and hours are slightly less cyclical. Overall, our results that both the real wages of hires in entry-level jobs and of job stayers declined in response to the Great Recession are robust to the specification of the second-step regression. The finding that hiring hours declined more than for job stayers is also robust, except for the

TABLE 4: Estimated semi-elasticity of real wages and hours with respect to the unemployment rate, 1998-2016: varying the specification of the second-step regression

	Wages		Hours	
	New hires (1)	Job stayers (2)	New hires (3)	Job stayers (4)
1. Baseline (OLS)	-2.83^{***}	-2.60^{**}	-1.47^{***}	-0.20
	(0.87)	(1.13)	(0.42)	(0.22)
2. Baseline with quadratic trend	-2.33^{***}	-1.96^{***}	-1.42^{***}	-0.13
	(0.48)	(0.51)	(0.46)	(0.15)
3. First differences (OLS)	-1.64^{***}	-1.84^{***}	-0.23	-0.09
	(0.59)	(0.40)	(0.72)	(0.16)
4. Baseline sample, but weighted by number of jobs per year	-2.64^{***}	-2.43^{**}	-1.49^{***}	-0.10
	(0.78)	(0.95)	(0.48)	(0.14)

Notes.- second-step regression results of estimated period effects on unemployment rate, $\hat{\gamma}$. First row is identical to Table 3, included here for comparison. Second row shows estimates when the second-step includes an additional quadratic time trend term. Third row estimates (2) in first differences, so measures the response of the log change in wages to a one percentage point increase in the change in unemployment. Fourth row applies WLS in the second step, with weights in proportion to the number of jobs observed per year.

Newey-West standard error estimates robust to first-order serial correlation in parentheses.

*** Statistically significant at the 1% level; ** at the 5% level, two-sided tests.

first-differenced version of (2), which indicates that the decrease in hiring hours is better understood as a medium-run and persistent development since 2008. In Appendix B we discuss the results of further robustness checks, which also do not affect our confidence in the main results.

4.2 Using labour productivity as the business cycle indicator

As an alternative indicator for the Great Recession we consider labour productivity, measured by log real gross value added per hour.[7] Measures of labour productivity are particularly relevant for a firm's hiring decisions. As Haefke et al. (2013) explain, the estimated response to this measure has an intuitive interpretation in standard search and matching models of the labour market: if real wages are perfectly rigid, then they should not respond to labour productivity, while a one-to-one response indicates fully

[7] Source: ONS Labour Market Statistics, April 2017, available at https://www.ons.gov.uk/.../apr2017; accessed 24/04/2017.

flexible wages.[8] Table 5 shows the estimated elasticity when we use labour productivity instead of the unemployment rate in regression (2). In the first row we use aggregate labour productivity as the business cycle indicator. The estimates are significantly smaller than one, but positive. Hiring hours also respond significantly, though less than real wages.

TABLE 5: Estimated elasticity of real wages and hours with respect to labour productivity, 1998-2016

	Wages		Hours	
	New hires (1)	Job stayers (2)	New hires (3)	Job stayers (4)
1. Labour productivity (I) whole economy	0.82^{***} (0.09)	0.94^{***} (0.07)	0.24^{**} (0.09)	-0.13^{***} (0.04)
2. Labour productivity (II) services sector	0.88^{***} (0.12)	1.02^{***} (0.09)	0.27^{**} (0.09)	-0.01 (0.04)

Notes.- second-step regressions of estimated period effects on alternative indicator of the business cycle. First-step estimated according to (1) and (3). "Labour productivity (I)" uses the log of real whole economy gross value added (GVA) per hour: ONS series LZVB. "Labour productivity (II)" uses the log of real gross value added (GVA) per hour in Services (sectors G-U): ONS series DJP9. We adjust both series by multiplying by the ratio of CPI to Producer Price Index of the services sector.
Newey-West standard error estimates robust to first-order serial correlation in parentheses.
*** Statistically significant at the 1% level; ** at the 5% level, two-sided tests.

Because over 90 per cent of jobs in our baseline sample belong to the services industry, and the response of labour productivity to the Great Recession was not the same across sectors, we also use labour productivity of the services sector as the cyclical indicator. The second row of Table 5 shows a higher estimated elasticity of real wages and hours worked with respect to services sector labour productivity. Real wages of new hires and job stayers significantly decrease by 0.9 and 0.8 per cent when aggregate labour productivity decreases by one per cent. The difference between these values is insignificant. The estimates for job stayers and new hires do not significantly differ from one at the 99 per cent and 95 per cent confidence level respectively: suggestively, real wages in the UK were perfectly flexible to aggregate labour productivity over the sample period.

[8]Appendix Figure D3 shows the time series of each business cycle indicator used and Appendix Table C5 shows the underlying values.

Our estimated hiring wage elasticity is of a comparable magnitude to that found by Haefke et al. (2013). These authors find an elasticity for the real wages of new hires of around 0.8 with respect to real output per hour in the non-farm business sector in the US. Similarly, Carneiro et al. (2012) find that the real wages of both stayers and hires increase approximately one-to-one with aggregate real output per worker in Portugal. Stüber (2017) finds that average real daily earnings of incumbent German workers increase by 0.5 per cent if aggregate real output per worker increases by one per cent, and he estimates a significantly smaller coefficient for new hires.

4.3 The role of the National Minimum Wage

Our results suggest that the real wages of new hires are just slightly more responsive to business cycle conditions than for job stayers. One potential explanation for this finding is the presence of a wage floor. This could constrain firms in how far they can reduce hiring wages. In 1999 such a floor was introduced in the UK in the form of the National Minimum Wage (NMW), with both adult and youth rates applying nationwide. These are usually uprated on an annual basis.[9] Collectively bargained wages can also limit a firm's flexibility in setting hiring wages. However, at the onset of the Great Recession, only six per cent of new hires in our sample were covered by a national or industry-level collective agreement (affecting working conditions, not necessarily pay). Therefore we consider the NMW to be the more likely limit on the responsiveness of hiring wages.

Figure 3 displays the real NMW rate that applied to workers aged 21 and older, along with the 10th percentile, 25th percentile, and median real wages of new hires within entry-level jobs for each year.[10] These hiring wages are not adjusted for changes in sample composition and include only workers aged 22-64. Between 2006 and 2015 new hires at the 10th percentile of the wage distribution were paid the legal minimum, i.e. the real value of the adult rate. In 2016, the 10th percentile of new hires increased more than the adult rate, which followed the introduction of a higher NMW rate for workers aged 25 and over. We also observe a narrowing of the gaps between the minimum wage

[9]Source: https://www.gov.uk/national-minimum-wage-rates; accessed 01/07/2017.
[10]The adult rate age limit was decreased from 22 to 21 in 2010.

and both the 25th percentile and median of hiring wages over the sample period. In other words, the domain of the distribution of real hiring wages at the job level, for employees aged 22-64, became more restricted from below at the level of the real NMW adult rate during the recent downturn. The wages of job stayers in CH-firms were less constrained by the minimum wage than hiring wages, since stayers are generally paid more than new hires (see Figure D1D).

FIGURE 3: Real hourly wages of new hires and NMW adult rate, ages 22-64

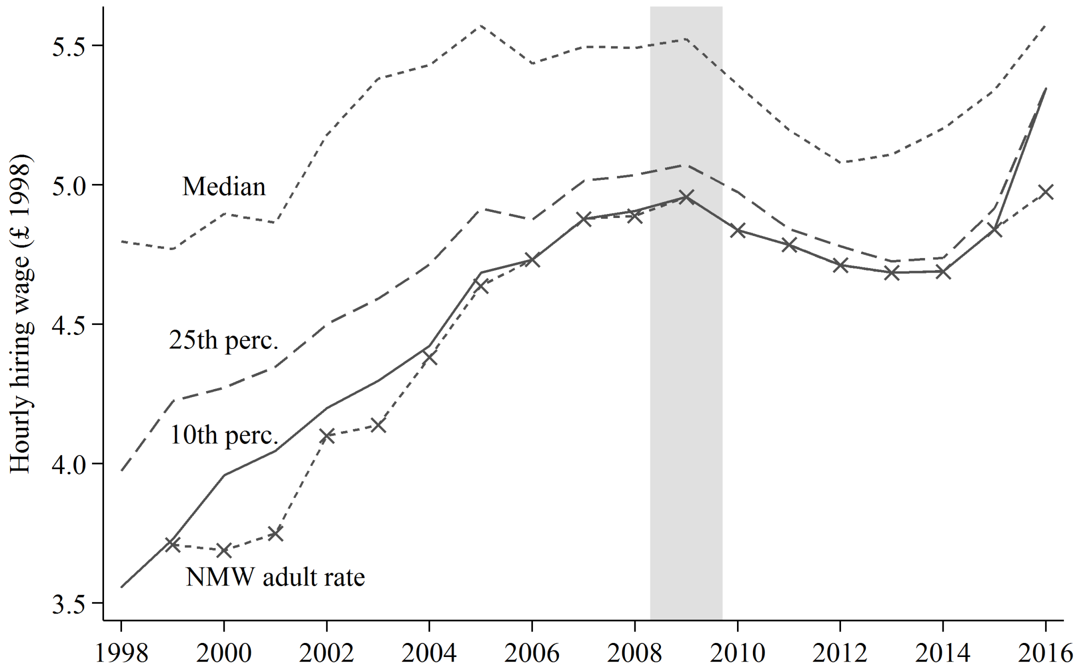

Notes.- National Minimum Wage adult rate and 10th, 25th, and median percentile of job-level hourly hiring wages, ages 22-64. All monetary values are deflated to 1998 values using the CPI. Shaded area marks official UK recession.

To answer the question of how hiring wages might have responded to the Great Recession in the absence of a binding minimum wage, we use the kernel re-weighting method of DiNardo et al. (1996). A description of this method is provided in Appendix F. A partial equilibrium assumption underlies this method: the number and composition of entry-level jobs is not affected by the NMW. This assumption is unlikely to hold in reality. Nevertheless, this method allows us to assess the impact of the NMW in a simple and transparent way.

Here we briefly explain the intuition. For each year following 2004, we replace the density of job-level real hiring wages which was at or below the real value of the NMW in that year, with the corresponding section of the 2004 density, adjusted for differences in observable job characteristics. Then, we re-scale this counterfactual density so that the two sections integrate to one. We select 2004 as the base year because this was the last year when the real value of the NMW was below its lowest level in 2014 (see Figure 3). For this estimation we use the plug-in method of Sheather and Jones (1991) to select the optimal bandwidth, which ranges from 0.01 to 0.04 for our sample.

The most important parameter in this kernel re-weighting exercise is the assumed size of the spillover effect of the minimum wage, i.e. the highest value of the real hiring wage density which is affected by the NMW. The more spillover we assume in a period, the more of this period's density - the section below the real value of the minimum wage plus any spillover - is replaced with the corresponding section of the 2004 density.

To the best of our knowledge, the extent to which the minimum wage is affecting real hiring wages in the UK has not been addressed. Varying estimates exist for the size of the spillover on UK wages in general, with estimates ranging from almost no spillover effects (Dickens and Manning, 2004) to relatively small effects up to the 5th percentile of wages above the NMW (Stewart, 2012), and up to 40 per cent above the NMW (Butcher et al., 2012). Therefore, we estimate counterfactual real wage densities for new hires, assuming spillover effects ranging from 0 to 10 p.p. above the real NMW in a given year.

To compute hiring wages at the job-level from counterfactual densities, we assume that the rank of a job in the distribution of hiring wages is preserved under different values of the NMW. Then we re-estimate regressions (1) and (2), using each of the counterfactual real hiring wage samples estimated with different spillover parameters. Figure 4 displays the point estimates of the counterfactual semi-elasticity of real hiring wages with respect to the unemployment rate across a range of assumed parameters of the spillover.

Assuming that there is no spillover effect, the left-most circle shows that the responsiveness of real hiring wages to the unemployment rate increases from -2.83 to -2.97 per cent. The standard errors are comparable to the baseline value (0.9)

FIGURE 4: Counterfactual estimates of the semi-elasticity of real hiring wages, 1998-2016: varying the assumed spillover effect of the NMW

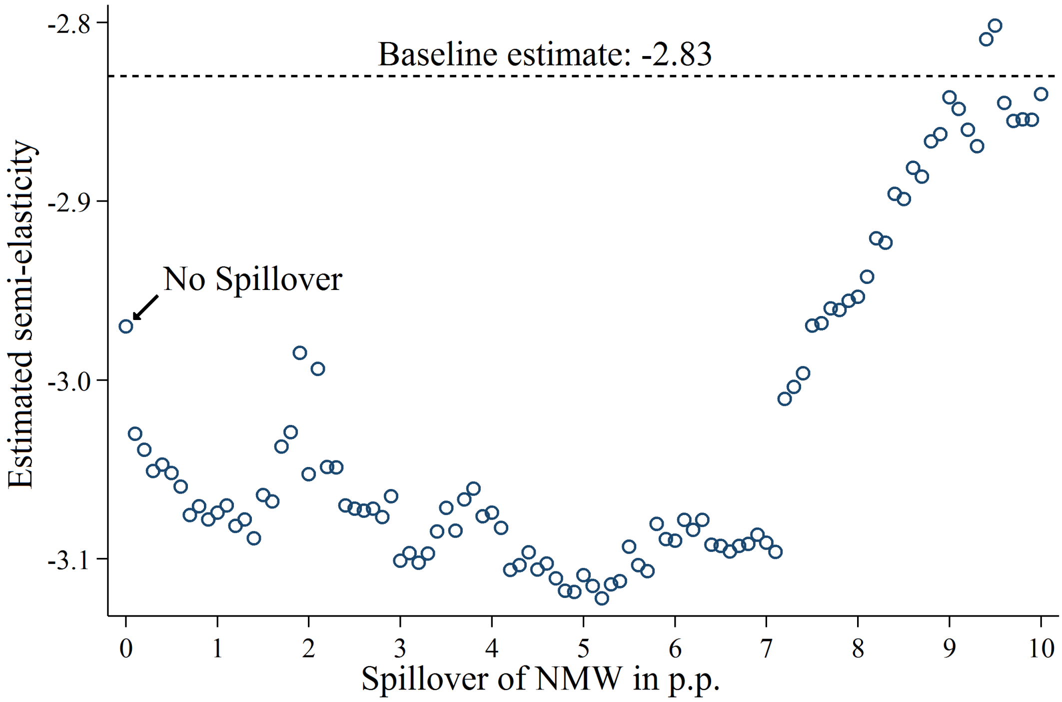

Notes.- each circle represents an estimate of the semi-elasticity of real hiring wages with respect to the unemployment rate. Standard errors lie outside of the figure. The horizontal axis shows the assumed spillover effect in p.p.. Dashed line shows the baseline estimate of semi-elasticity. We use a Gaussian kernel, and the bandwidth is selected using the Sheather-Jones plug-in estimator.

and lie outside the range of this figure. The semi-elasticity falls below -3.1 when the spillover effect increases to five p.p. If the spillover increases above seven p.p., then the responsiveness of real hiring wages to changes in the unemployment rate begins to decrease back towards the baseline estimate: the shape of the counterfactual density increasingly resembles the shape of the density observed in 2004 when we assume larger spillover effects, and hence the variation of hiring wages over time in entry-level jobs declines. These results suggest that the NMW constrained firm in how far they could reduce wages of new hires during the Great Recession.

5 How did wages and hours evolve after hiring?

FIGURE 5: Paths of real wages and hours for cohorts of new hires

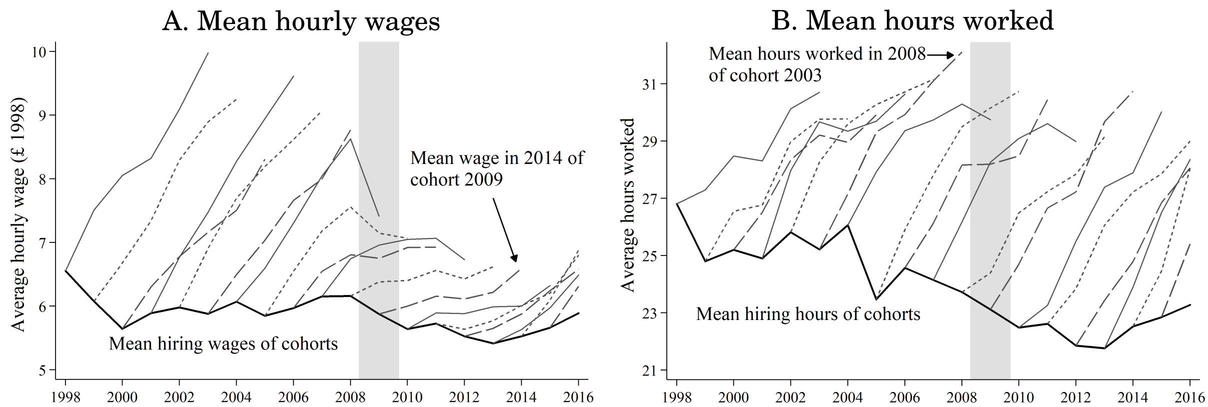

Notes.- the solid lines give the average real hiring wage and weekly hours worked for each cohort of new hires in our sample of entry-level jobs (i.e. column (1), Table 1). Each line branching off from the solid line shows the paths of wages or hours of these hiring cohorts over time, as their tenure in the job increases. When employees leave their hiring jobs they also exit the samples of their respective cohorts.

So far we have demonstrated that both real hiring wages and hours worked in entry-level jobs significantly decreased during the UK's Great Recession. However, Haefke et al. (2013) and Elsby et al. (2016) argue that a firm's decision to hire an additional worker should depend on the expected present value of the marginal profit from a successful match. The initial hiring wage and hours worked only form part of this expected value, with hours only relevant if there are non-linearities in the firm's production or labour cost functions. If firms who can hire at lower wages and hours during a recession also have to deliver greater wage growth in the job, then the expected present value of the marginal product is potentially less cyclical than measured for the hiring wage. Thus our previous estimates of wage flexibility may be less important for understanding the muted employment response of the UK's Great Recession than first imagined.

As an initial assessment of the importance of cohort effects, Figure 5 plots the real hourly wages and hours worked averaged over employees instead of jobs, for each cohort of entry-level new hires, conditional on these employees staying in their respective jobs.

The average hiring wages and hours in each year are shown as solid lines. Panel A suggests that wages exhibit cohort effects: the real wages of hiring cohorts from 1998 to 2005 mostly seem to have parallel trends in the first three years on the job, similar to the findings of Baker et al. (1994) for one US firm. But, unlike these authors, we see that the cohort-specific paths of wages respond to the business cycle, as shown by a decline in wage growth during the years of the Great Recession. Cohorts hired during this time seem to be locked into low wage growth trajectories. For example, the mean wages of the 2013 cohort in 2015 were still below the mean wages of the 2014 cohort in 2015. Panel B of Figure 5 similarly suggests that the path of hours worked depends on cohort effects, though less strikingly so than for wages, as growth trends remained mostly parallel throughout the period. In other words, differences in cohort hiring wages and hours over the business cycle seem to persist, and may even reinforce the initial decline in labour costs.

Comparing sample averages over time is likely to be subject to a composition bias, since relatively low-wage employees are given less weight during downturns than in normal times (Solon et al., 1994). Therefore we estimate how the wages and hours of new hires in entry-level jobs evolved over three years of subsequent tenure, including match-fixed effects to control for the changing composition of matches over the business cycle. We include only consecutive observations of a worker in some job, such that a worker with three years of tenure must be observed the previous two years. The sample of workers for each hiring year is unbalanced, since workers exit from entry-level jobs: either they switch jobs within the same firm or across firms, or they exit into non-employment. Using least squares we estimate

$$w_{m\tau} = \theta_m + \psi_{S(m)\tau} + \boldsymbol{x}'_{m\tau}\boldsymbol{\phi} + \eta_{m\tau} \;, \tag{4}$$

where the dependent variable is the log real wage of some match m between a worker (i) and job (j) with tenure τ, where $S(m)$ is a function indicating in which year s a match was formed. θ_m is a match-fixed effect and $\psi_{s\tau}$ are cohort-tenure-fixed effects for any matches beginning in year $s = 2001, \ldots, 2013$ with years of tenure $\tau \in [0,3]$. The sample size per cohort is initially over a thousand employees with tenure greater than a year, and then declines to around four hundred employees per cohort with tenure over three years. The vector $\boldsymbol{x}_{m\tau}$ contains time-varying quadratic controls for the size of the firm, and $\eta_{m\tau}$ is the error term. We estimate (4) by excluding the effects ψ_{s0}, so the estimated values of $\hat{\psi}_{s\tau}$ for $\tau > 0$ are interpreted as log changes relative to the hiring wages in entry-level jobs. Although there is certainly endogenous selection of the employees who stay in these jobs for up to three years, the match-fixed effect should partially address this concern. We similarly estimate (4) with log basic weekly hours worked as the dependent variable.

The sample of new hires used for this analysis is a subset of our baseline sample, because we require at least one completed year of tenure. Therefore, some jobs included in the baseline sample are no longer represented here: the sample size of jobs is around 25 per cent smaller. The estimated real wage semi-elasticity of new hires with respect to the unemployment rate for this group of workers is 3.2, and is slightly larger in absolute terms than in our baseline sample, while hours worked are just as responsive as measured before.

FIGURE 6: Estimated composition-adjusted and cohort-specific log changes $\hat{\psi}$ in real wages and hours relative to hiring levels: workers who stay in entry-level jobs

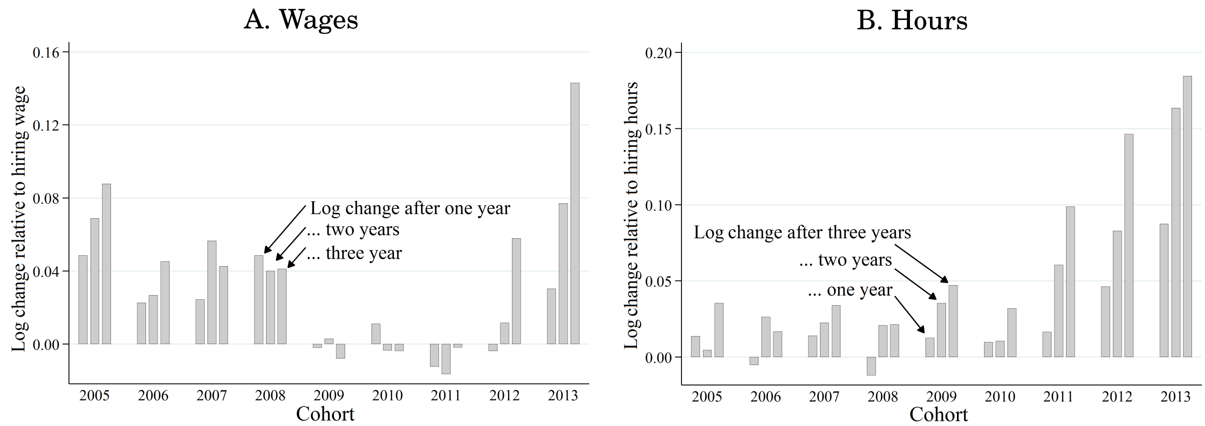

Notes.- cohort average change in log wages and log hours with tenure, relative to respective hiring values, in entry-level jobs. Composition adjusted by controlling for match-fixed effects. See Appendix Table C6 for standard errors and results for all other hiring years in 2001-13.

We plot the estimated cohort-tenure-fixed effects in Figure 6 for selected hiring years, and the underlying estimates are displayed in Appendix Table C6 for all years with confidence intervals. The last cohort of hires unaffected by the Great Recession within three completed years of tenure was 2005. Panel A shows a clear U-shaped response of real wage growth at all levels of tenure over the Great Recession. Over the first year on the job, the wages of workers hired in 2009 were stagnant, while for those hired before the recession they grew on average by two percent, and for those in the last cohort by three per cent. Similarly, there was no real wage growth over the subsequent three years for 2009 hires, compared with over eight per cent for 2005 hires.

The cyclical differences in hours growth in these jobs are less pronounced. Initially the increases over three years were smaller in 2008 than pre-recession, with negative growth over the first year of tenure. However the average increases following the recession were greater. This latter period coincides with the persistent rise and peak in part-time employment in the UK following the financial crisis. In roughly equal parts, the changes of hours worked within jobs, represented by the data here, are due to switches between part- and full-time work and increasing hours within these categories. Thus the pattern within these particular jobs suggests a caveat to the findings of Borowczyk-Martins and Lalé (2017), who use worker-level flows to show that within employment switches mostly accounted for the rise and persistence of part-time employment during the UK's Great Recession. Their finding of cyclical transition rates at the average worker level, which go against our job-level results, could be the outcome of job switching within the same firm. Also, the findings of Borowczyk-Martins and Lalé could mostly apply to workers with longer tenure than three years or shorter than one year. Similarly, Kurman and McEntarfer (2017) document that employees who stay for at least two years in the same firm, as opposed to the same job, experience cyclical variation in hours worked. Future research should try to address whether or not a large part of the measured worker-level cyclical hours adjustments at the average (or aggregate) level involves cyclical job switching, if not also firm switching.

The findings in this section suggest that firms were not only able to significantly reduce the real wages and hours of new hires in response to the Great Recession, but also depress wage growth with subsequent tenure. However there are at least two reasons this evidence is only suggestive. First, it only applies to workers who stay in the exact same job in the firm, whereas in reality, expected employee progression or reallocation to other jobs within the firm also affects the ex ante present value of a match and the hiring decision. Second, the regression in (4) is subject to the same measurement criticism which the majority of this paper shows is important: it does not control for the endogenous selection and weighting of matches over time, which we are unable to adequately address due to a small number of degrees of freedom at the job level here.

6 Discussion of findings and possible explanations

Using essentially the same dataset but without firm-identifiers, Elsby et al. (2016) show that UK real wages behaved very differently during the Great Recession when compared with previous recessions: during the 1980s and 1990s downturns the growth in real wages for British job stayers slowed, whereas it turned markedly negative in the most recent downturn. This matches the findings of Gregg et al. (2014), who document that UK wages became significantly more sensitive to changes in local unemployment rates sometime in the early 2000s. Both Elsby et al. (2016) and Gregg et al. (2014) emphasise that the decline in unionisation in the UK since the 1970s could only account for a small part of these observed changes in the behaviour of real wages. One argument for this is that US real wages remained relatively constant in the years following the 2008 financial crisis, while US employment fell sharply, despite the US seeing a greater decline and lower contemporary level of unionisation than the UK.

A similar line of argument applies for the role inflation. In both countries price inflation was historically low before and during the Great Recession. In Appendix E we further dispel the notion that price inflation could account for the high level of real wage flexibility in the Great Recession, by demonstrating that there is a lack of absolute nominal wage rigidity among UK job stayers. We extend the time period of Elsby et al.'s account of UK nominal wage rigidity, and specifically consider year-to-year hourly wage changes among the job stayers in the baseline sample from our main analysis. As many as two-thirds of these employees experienced annual real wage cuts at the height of the downturn, while around a quarter also experienced nominal wage cuts. The incidence of exactly zero annual nominal wage changes increased from approximately 0-2 per cent of employees before 2008-09 to 3-5 per cent in the years after. Our main findings on the extent of UK real wage flexibility reflect the fact that large numbers of employees experience yearly nominal wage cuts, almost independently of the economic cycle.

Blundell et al. (2014) argue that the UK's labour supply curve shifted to the right during the Great Recession. This was most likely caused by welfare reforms, which led to the addition, and stricter enforcement, of job search requirements for several groups of non-employed persons. For example, lone parents, who constitute approximately a quarter of all UK family households, were particularly affected. The age of the youngest child, at which lone parents are entitled to unconditional income support, was gradually reduced from sixteen to five years old between 2008 and 2012. If their youngest child was older than these lowered thresholds, then lone parents would have had to show evidence that they were searching for work in order to receive the same income support as they were entitled to previously without searching. It has been estimated that these particular policy changes led to an increase of almost ten per cent in the employment rate among UK lone parents, despite this occurring throughout a major recession (Avram et al., 2016). It is plausible that increased competition for jobs, brought on by the cumulative and extensive changes in the UK's active labour market policy since the last major downturn in the early 1990s, resulted in large decreases in the real values of workers' reservation wages and outside options, and thus led to new hires and job stayers accepting large decreases in real wages.

Perhaps our most striking finding for the behaviour of the UK labour market since 2008 is the extent to which hiring hours in jobs were reduced. How could this shift from full- to part-time recruitment be explained? Shifts in the labour supply curve, particularly for part-time work, are again potentially relevant. The UK has a system of tax credit benefits for working families with children similar to the US earned income tax credits. Entitlement for the work-contingent component requires at least one adult to work for a minimum of sixteen hours per week. There is observable bunching in the distribution of employee hours worked around the thresholds in the UK tax credits system, which is unsurprising given the large differences in the amount of credits families receive around these levels (see Blundell et al. (2016) for a more detailed discussion). This part of the UK welfare system cushions workers from income loss when their working hours decline, as well as encouraging them to take part-time work more readily than they perhaps would otherwise. In fact, the number of people in the UK who said that they were working part-time because they could not find a full-time job in 2013 stood at the highest level on records: almost 1.5 million (6 per cent of all employees), compared with 2.5 million unemployed, and compared with 0.7 million involuntary part-time employed in 2007.[11]

Another possible cyclical feature of labour markets is the so-called "Added Worker Effect", whereby individual household members will increase their labour supply when the household experiences persistent income shocks, typically thought of as resulting from a partner's job loss. There is some aggregate evidence of this effect for the UK, based on individual-level labour force transition rate data (Razzu and Singleton, 2016). However, Bryan and Longhi (2013) have shown that while this effect seems to draw individuals into the UK unemployment pool, it does not significantly increase their likelihood of becoming employed. The added worker effect is therefore unlikely to be a large part of the overall story of why hiring hours were flexible since 2008.

Montgomery (1988) discusses the factors which determine firms' demand for part-time employees. If there are fixed costs of hiring and training new employees, then these costs are unlikely to vary between part- and full-time hires in the same job: the ratio of hours to fixed costs will often be lower for part-time hires. Firms require compensation for this lower return from part-time hiring, that is, the hourly wage per worker has to be lower. This firm-side compensating differential should be stronger for higher-skilled jobs, where hiring and training costs are typically greater. Montgomery (1988) provides evidence for these features of wage-setting and hiring behaviour in the presence of fixed costs among US establishments. Moreover, if firms have to pay

[11] Source: ONS Labour Market Statistics, October 2017, available at https://www.ons.gov.uk/.../october2017; accessed 07/11/2017. See also Bell and Blanchflower (2013) for more details about the so-called "Underemployment" in the UK.

all workers in some job the same hourly rate, then firms are more likely to employ full-time employees when there are fixed hiring costs. However, fringe benefits (pension contribution, health care) function as quasi-fixed costs which might only be offered to full-time employees, and thus shift the demand from full- to part-time workers. To the extent that these fixed costs depend on the level of productivity, it is possible that they decline during recessions, and thus make part-time hiring more likely. The cyclical properties of fixed hiring costs in the UK is an interesting empirical question for future research.

In summary, some combination of increasing labour supply and the institutional framework surrounding the UK's labour market are the most likely explanations of our main findings. However more research is needed to understand if this flexibility over the business cycle in working conditions will become the new normal for the UK labour market. Further, the novel fact documented here regarding the extent of hours reductions in job-level hires over the business cycle should be explored outside the specific context of the UK's Great Recession.

7 Conclusion

We provide new estimates on the flexibility of UK wages during the Great Recession. Most importantly this is measured at the job level, which is the correct approach to understanding how firms adjust their labour costs in response to business cycle conditions in frictional labour markets. We find that job-stayer real wages respond by as much as 2.6 per cent for every one percentage point rise in the unemployment rate. Their elasticity with respect to aggregate labour productivity equals approximately one. Hiring wages are at least as responsive to the business cycle as the wages of job stayers. This conforms with results from other countries, suggesting that rigid hiring wages are not the appropriate way to model and understand the observed fluctuations in unemployment.

Several other studies have also measured real wages in Britain's Great Recession, concluding that the magnitude of their response likely explains the high-employment and low-productivity experience of the subsequent decade, compared with previous downturns and other countries (Blundell et al., 2014; Gregg et al., 2014; Elsby et al., 2016). Once we strip away cyclical job composition bias, our estimates of the real wage response are a magnitude greater than found in these previous studies. While this large and significant wage response now seems even more likely to account for the UK economy's unusual experience of the Great Recession, the puzzle still remains as to why firms were able to adjust wages so freely, and why workers were so willing to accept these changes.

To the best of our knowledge, this is the first paper to combine the robust job-level measurement of cyclical responses in real wages with hours worked for new hires and job stayers, within the same methodological framework. We find that the hours worked by job stayers did not respond to the Great Recession. Conversely, the hours of new hires among the same firms responded significantly, decreasing by 1.5 per cent for every one percentage point rise in the unemployment rate, mostly through firms switching between full- and part-time workers. We believe this is a new empirical account of cyclical firm behaviour, which should in the first instance be tested outside the specific UK context, and subsequently reflected on when modelling how firms adjust their workforces to shocks.

We also find evidence that hours response estimates, like wages, can be subject to a large bias induced by the endogenous cyclical selection of jobs, though this is pro- as opposed to countercyclical as in the case of wages. Some recent studies have explained procyclical average hours worked in the whole economy by changes in worker transition rates between part- and full-time employment. However, changes in aggregate hours, like real wages, tell us very little about what happens at the job level, where we find no significant response to the unemployment rate for employees who stayed with the same job and firm. The robust distinction here between the responses of wages and hours within jobs also offers insight into the results for Germany in Stüber (2017), which are somewhat atypical for this literature: wages in these German data are perhaps better interpreted as average daily earnings.

While the approach to measurement here is inspired by Solon et al. (1994) and closely follows Martins et al. (2012), by forgoing representativeness for greater certainty on what wage responses are actually being identified, we also offer some original methodological insights. Unlike previous job-level studies, we are sure to compare the wages of new hires and job stayers within the same sample of firms, and in the latter case also account for endogenous cyclical selection. This enables us to be more confident when comparing the estimated hiring and job-stayer responses to the business cycle.

We also offer some evidence that the UK's National Minimum Wage restricts how far firms can reduce wages, and our estimate of hiring wage flexibility could have been even greater without this restraint. In this regard, it is surprising that other related studies do not similarly consider this when interpreting their main findings, given that elsewhere and historically large fractions of employees and jobs could be subject to tight and infrequently negotiated (collectively bargained) wage floors.

Cohorts hired during the Great Recession were not only paid lower wages initially, but were also locked into low-wage growth paths. This significantly reduced the present value of labour costs from the firm's perspective for hires made during this time. In this respect, it seems that firms' hiring wages were even more flexible than our results for the initial real wages of new hires show. We therefore take our results as evidence against any theory that hiring wages are especially rigid. Moreover, when combined with the shift from full- to part-time hiring, firms were able to significantly reduce their labour costs per new employee.

References

Abraham, K. G., and J. C. Haltiwanger. 1995. "Real Wages and the Business Cycle." *Journal of Economic Literature*, 33(3): 1215–1264. [page 2]

Angrist, J. D., and J.-S. Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press. [page 6]

Avram, S., M. Brewer, and A. Salvatori. 2016. "Can't work or won't work: Quasi-experimental evidence on work search requirements for single parents." IFS Working Papers W16/11, Institute for Fiscal Studies. [page 31]

Baker, G., M. Gibbs, and B. Holmstrom. 1994. "The Wage Policy of a Firm." *The Quarterly Journal of Economics*, 109(4): 921–955. [page 27]

Bell, D. N., and D. G. Blanchflower. 2013. "Underemployment in the UK Revisited." *National Institute Economic Review*, 224(1): F8–F22. [page 32]

Blundell, R., M. Brewer, and M. Francesconi. 2008. "Job Changes and Hours Changes: Understanding the Path of Labor Supply Adjustment." *Journal of Labor Economics*, 26(3): 421–453. [page 15]

Blundell, R., M. Costa Dias, C. Meghir, and J. Shaw. 2016. "Female Labor Supply, Human Capital, and Welfare Reform." *Econometrica*, 84(5): 1705–1753. [page 32]

Blundell, R., C. Crawford, and W. Jin. 2014. "What Can Wages and Employment Tell Us about the UK's Productivity Puzzle?" *Economic Journal*, 124(576): 377–407. [page 1, 2, 15, 31, 34]

Borowczyk-Martins, D., and E. Lalé. 2017. "Employment Adjustment and Part-time Jobs: The U.S. and the U.K. in the Great Recession." IZA Discussion Papers 9847, Institute for the Study of Labor (IZA). [page 15, 29]

Bryan, M., and S. Longhi. 2013. "Couples' Labour Supply Responses to Job Loss: Boom and Recession Compared." IZA Discussion Papers 7775, Institute for the Study of Labor (IZA). [page 32]

Butcher, T., R. Dickens, and A. Manning. 2012. "Minimum wages and wage inequality: Some theory and an application to the UK." Discussion Paper 1177, Centre for Economic Performance. [page 24]

Carneiro, A., P. Guimarães, and P. Portugal. 2012. "Real Wages and the Business Cycle: Accounting for Worker, Firm, and Job Title Heterogeneity." *American Economic Journal: Macroeconomics*, 4(2): 133–152. [page 4, 22]

Crawford, C., W. Jin, and H. Simpson. 2013. "Productivity, Investment and Profits during the Great Recession: Evidence from UK Firms and Workers." *Fiscal Studies*, 34(2): 153–177. [page 1]

Devereux, P. J. 2001. "The Cyclicality of Real Wages within Employer-Employee Matches." *Industrial and Labor Relations Review*, 54(4): 835–850. [page 3]

Devereux, P. J., and R. A. Hart. 2006. “Real Wage Cyclicality of Job Stayers, Within-Company Job Movers, and Between-Company Job Movers.” *Industrial and Labor Relations Review*, 60(1): 105–119. [page 3, 19]

Dickens, R., and A. Manning. 2004. “Spikes and spill-overs: The impact of the national minimum wage on the wage distribution in a low-wage sector.” *Economic Journal*, 114(494): 95–101. [page 24]

DiNardo, J., N. M. Fortin, and T. Lemieux. 1996. “Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach.” *Econometrica*, 64(5): 1001–1044. [page 4, 23, 55, 56, 58]

Donald, S. G., and K. Lang. 2007. “Inference with Difference-in-Differences and Other Panel Data.” *The Review of Economics and Statistics*, 89(2): 221–233. [page 6]

Elsby, M., D. Shin, and G. Solon. 2016. “Wage Adjustment in the Great Recession and Other Downturns: Evidence from the United States and Great Britain.” *Journal of Labor Economics*, 34(S1): S249 – S291. [page 3, 14, 26, 30, 34, 53]

Gertler, M., C. Huckfeldt, and A. Trigari. 2016. “Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires.” CEPR Discussion Papers 11318, C.E.P.R. Discussion Papers. [page 3]

Gertler, M., and A. Trigari. 2009. “Unemployment Fluctuations with Staggered Nash Wage Bargaining.” *Journal of Political Economy*, 117(1): 38–86. [page 3, 18]

Gregg, P., S. Machin, and M. Fernández-Salgado. 2014. “Real Wages and Unemployment in the Big Squeeze.” *Economic Journal*, 124(576): 408–432. [page 1, 30, 34]

Haefke, C., M. Sonntag, and T. van Rens. 2013. “Wage rigidity and job creation.” *Journal of Monetary Economics*, 60(8): 887–899. [page 20, 22, 26]

Hart, R. A., and J. E. Roberts. 2011. “Job Re-grading, Real Wages, and the Cycle.” IZA Discussion Papers 5912, Institute for the Study of Labor (IZA). [page 3]

Jäger, K. 2016. “EU KLEMS Growth and Productivity Accounts 2016 release - Description of Methodology and General Notes.” [page 12]

Kurman, A., and E. McEntarfer. 2017. “Downward Wage Rigidity in the United States: New Evidence from Administrative Data.” mimeo. [page 29]

Martins, P. S., G. Solon, and J. P. Thomas. 2012. “Measuring What Employers Do about Entry Wages over the Business Cycle: A New Approach.” *American Economic Journal: Macroeconomics*, 4(4): 36–55. [page 3, 4, 5, 9, 10, 15, 17, 18, 19, 35]

Montgomery, M. 1988. “On the Determinants of Employer Demand for Part-Time Workers.” *The Review of Economics and Statistics*, 70(1): 112–117. [page 32]

Nickell, S., and G. Quintini. 2003. “Nominal Wage Rigidity and the Rate of Inflation.” *Economic Journal*, 113(490): 762–781. [page 7, 53, 54]

Pessoa, J. P., and J. Van Reenen. 2014. “The UK Productivity and Jobs Puzzle: Does the Answer Lie in Wage Flexibility?” *Economic Journal*, 124(576): 433–452. [page 15]

Pissarides, C. A. 2009. "The unemployment volatility puzzle: Is wage stickiness the answer?" *Econometrica*, 77(5): 1339–1369. [page 3]

Razzu, G., and C. Singleton. 2016. "Gender and the business cycle: An analysis of labour markets in the US and UK." *Journal of Macroeconomics*, 47(Part B): 131–146. [page 32]

Schaefer, D., and C. Singleton. 2017. "Recent Changes in British Wage Inequality: Evidence from Firms and Occupations." 2017 Meeting Papers 459, Society for Economic Dynamics. [page 8]

Sheather, S. J., and M. C. Jones. 1991. "A Reliable Data-Based Bandwidth Selection method for Kernel Density Estimation." *Journal of the Royal Statistical Society. Series B (Methodological)*, 53(3): 683–690. [page 24]

Shin, D. 1994. "Cyclicality of Real Wages among Young Men." *Economics Letters*, 46(2): 137–142. [page 3]

Solon, G., R. Barsky, and J. A. Parker. 1994. "Measuring the Cyclicality of Real Wages: How Important is Composition Bias?" *The Quarterly Journal of Economics*, 109(1): 1–25. [page 2, 5, 27, 35]

Solon, G., S. J. Haider, and J. M. Wooldridge. 2015. "What Are We Weighting For?" *Journal of Human Resources*, 50(2): 301–316. [page 15]

Stewart, M. B. 2012. "Quantile Estimates of Counterfactual Distribution Shifts and the Effect of Minimum Wage Increases on the Wage Distribution." *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 175(1): 263–287. [page 24]

Stüber, H. 2017. "The Real Wage Cyclicality of Newly Hired and Incumbent Workers in Germany." *Economic Journal*, 127(600): 522–546. [page 4, 18, 22, 34]

Yagan, D. 2017. "Employment Hysteresis from the Great Recession." NBER Working Papers 23844, National Bureau of Economic Research, Inc. [page 5]

Appendix A. Further description of the data and sample

In what follows we give some additional details regarding the datasets used and how we have constructed sub-samples thereof. All of the relevant documentation and variable descriptions attached to these datasets are publicly available from the UK Data Service. The ONS has also published various documents concerning the data quality and consistency of the ASHE. We will publish our replication files for the analysis and sample construction.

We focus on methodological details through the period 1998-2016. Throughout this period, the ASHE should be a true random sample of all employees in employment, irrespective of employment status, occupation, size of employer etc. Given the legal obligation of employers to respond using payrolls, it has a high response rate and is believed to be accurate. There is no cumulative attrition from the panel, as any individual not included in the ASHE in any year, for whatever reason, remains in the sampling frame the following year. Conditional on a hundred per cent response, the ASHE is a true one per cent random sample of employees: all with a National Insurance Number which has a numerical part ending in 14. However there are two major sources of under-sampling, both occurring if individuals do not have a current tax record. This could happen for some individuals who have very recently moved job, or for those who earn very little (mostly part-time) and are not paying income tax or National Insurance in the period when their employers are looked up. From 2004 the ASHE aimed to sample some of those employees under-represented. It added supplementary responses for those without a PAYE reference, and also attempted to represent employees whose jobs changed between the determination of the sampling frame in January and the reference period in April. Since the ONS states that the biases that these amendments were introduced to address were actually small, we do not believe they could affect our results substantially. The ASHE also introduced some imputations, using similar matched 'donor' observations where responses were, for example, missing an entry of basic hours but had recorded pay. These imputations were added for weighting purposes, but throughout our analysis we ignore the weights in the ASHE, since they are designed to make the aggregate results population representative in terms of worker observables, and are not firm-level.

From 2005 a new questionnaire was also introduced, which was intended to reduce the latitude for respondents' own interpretations of what was being asked of them. From 2007 there were further notable changes. Before occupations were classified as follows: if the respondent stated an employee's job had not changed in the past year the previous year's occupational classification was applied - otherwise, it was manually coded. Afterwards an automatic coding, text recognition, tool was used. "The effect of using ACTR was to code more jobs into higher paying occupations. The jobs that tended to be recoded into these higher paying occupations generally had lower levels of pay than the jobs already coded to those occupations. Conversely, they tended to have higher levels of pay than the other jobs in the occupations that they were recoded out of. The impact of this was to lower the average pay of both the occupation group that they had moved from and that they had moved to." From 2007 the sample size of the ASHE was reduced by 20 per cent, with reductions targeted at those industries exhibiting the least variation in earnings patterns.

We use the ASHE annual cross-sections for each year from 1998 to 2016 and construct a panel as follows: first, we merge the two separate cross-sections for the year 2010, where one contains occupations coded in SOC2010 and the other in SOC2000. This is done to match occupations across classification schemes for the same individuals. In case of multiple jobs per individual, we exclude non-main jobs. In case of missing main job markers, we impute these based on the job with the highest working hours. In a next step we link employees across consecutive years based on their unique identifiers. This enables us to impute missing enterprise reference numbers (entrefs) backwards, since the ASHE contains a variable which indicates whether an employee is holding the same job as in the last reference period. Note that this "same job" variable alone does not allow between-firm and within-firm job changers to be distinguished. Subsequent to linking two consecutive years in this way, we use local unit identifiers to impute missing entrefs across individuals within the same year (the ONS states that the local unit identifiers are not consistent across years, rather they are created to identify establishments within years). We continue to update missing entrefs in this way back to and including 1998. The number of observations with non-missing entrefs after imputation declines rapidly as we go further back in time. While for the years 2003-2016 we are only adding a couple of missing entrefs per year, prior to 2003, and especially prior to 2000, we are imputing almost all entrefs. We could also impute entrefs for 1997, but this year does not include the marker that indicates whether an individual is working in the same job, which is vital to our sample selection strategy.

We keep only observations for individuals aged 16-64, and which have not been marked as having incurred a loss of pay in the reference period through absence, employment starting in the period, or short-time working, and which are marked as being on an adult rate of pay (i.e. dropping trainees and apprenticeships). This is practically the same filter applied for annual ONS published results on UK "Patterns of Pay" using the ASHE. We drop observations with missing basic hours, gross weekly earnings, or hourly wage rates. Basic hours are intended to be a record for an employee in a normal week, excluding overtime and meal breaks. Gross weekly pay is the main recorded value in the survey, and from this overtime records are subtracted. Hourly rates are then derived from dividing by basic hours worked. We drop observations with over a hundred or less than one basic hour worked, as these could reflect measurement error and the inclusion of overtime. Full-time is defined as working over thirty basic hours in a week. But there are a tiny number of discrepancies in some years, we believe relating to teaching contracts, where the definition applied by the ONS differs. We however recode these such that for all observations the thirty hours threshold applies. To further address some potential for measurement error, especially in the recorded basic hours, we drop observations whose derived hourly rate of pay, excluding overtime, is less than 80 per cent of the applicable National Minimum Wage (NMW) each April, with allowance for the different age-dependent rates of the NMW over time. We set the threshold lower to avoid dropping observations where employers have rounded figures about the NMW, where the degree of rounding could vary with the actual value of the NMW, a behaviour which has been hypothesised by the ONS.

We define an entry or new hire into a firm as an individual with less than one year of tenure. For this we make use of the employment start date. The ASHE contains information on when an employee starting working for an enterprise from 2002 onwards. We drop a tiny number of unrealistic entry dates, where the start date lies either in the future or implies an employee started work aged fifteen or younger. Unfortunately there are some inconsistencies across years in these records. First, an employee can be employed by the same company for three consecutive years, holding the same job, but the starting dates recorded in the first and third years, though identical, can vary from the second. In this case we update the "one-off" deviation with the value of the previous year. Second, if we observe an employee in a chain of consecutive years in the same firm, holding the same job, but the start date differs for some years, then we impute the earliest date available. This decision is based on a conservative interpretation of a "new hire": in case of previous employment within the same firm, we do not include an employee in our CH-firms sample of new hires if we are in any doubt. Given our finding is that hiring earnings cyclicality is larger in absolute terms than that of job stayers, any expected bias would go in the opposite direction. Finally, we use employment start date to impute entrefs for employees backwards again. This enables us to no longer have to observe employees in a chain of consecutive years to make imputations. We then again use within-year local unit identifiers to update longitudinal entrefs within a year for other employees with missing entrefs. The ASHE contains the number of employees of an enterprise as listed in the Inter-Departmental Business Register (IDBR). A very small fraction of employees in the same enterprise and year have missing or varying values for this variable. We impute the same value for all employees within year and enterprise as the modal value for the firm.

For 1996-2001 occupations are classified using the three-digit ONS1990 Standard Occupational Classification (SOC). For 1998-2010 occupations are classified using the four-digit SOC2000, and for 2011-2016 with the SOC2010. We experimented using the ONS' publicly available cross-walk from 2010 and 2000, but discovered that this causes a large structural break in the distribution of occupations. In particular it causes a substantial additional degree of polarisation of work from 2002 onwards. Therefore we use our own cross-walk obtained from the ASHE cross-section 2010, as discussed above, to map SOC2010 into SOC2000 *within* an enterprise. However, some occupations for some firms are not observed in the year 2010, but are in the following years, for which we do not have double coded data. To address this we first convert SOC2010 to the 2008 International Standard Classification of Occupations (ISCO), obtained from the ONS website. Then we convert SOC2000 to ISCO1988, where we obtain conversion tables from the Cambridge Social Interaction and Stratification Scale (CAMSIS) project. Finally, we use the ISCO2008 to ISCO1988 cross-walk, available from the International Labour Organization. For the industry classification, we convert ONS Standard Industrial Classification (SIC) 2007 to 2003, using files made available by the UK Data Service. This conversion uses the 2008 Annual Respondents Dataset, where both classifications were applied, and where any 2007 code mapping to multiple 2003 codes is decided using whichever of the two bore a greater share of economic output.

Appendix B. Further robustness checks

Table B1 presents some further robustness checks of the main empirical results presented in Table 3. The first row repeats our baseline/main result for convenience. The robustness discussed here is with regards to the specification of the first step of the regression model: (1) & (3). The specification of the second step is unchanged compared with the baseline. The second row describes the estimated semi-elasticity of real wages and hours with respect to the unemployment rate when typical job-level measures are employee sample means, rather than median values. Qualitatively the results are unchanged: wages for hires and job stayers exhibit a sizeable and significant cyclical response, as do hiring hours, though the difference in wage response between hires and

TABLE B1: Estimated semi-elasticity of real wages and hours with respect to the unemployment rate, 1998-2016: more robustness checks

	Wages		Hours	
	New hires (1)	Job stayers (2)	New hires (3)	Job stayers (4)
1. Baseline	-2.83^{***}	-2.60^{**}	-1.47^{***}	-0.20
	(0.87)	(1.13)	(0.42)	(0.22)
2. Job means	-2.98^{***}	-2.67^{**}	-1.08^{***}	-0.28
	(0.91)	(1.17)	(0.31)	(0.18)
3. Baseline, but without controls	-2.78^{***}	-2.61^{**}	-1.49^{***}	0.02
	(0.80)	(1.10)	(0.48)	(0.21)
4. Baseline, but including public sector	-2.63^{***}	-2.62^{**}	-2.17^{***}	-0.41^{***}
	(0.95)	(1.31)	(0.30)	(0.10)
5. RPI instead of CPI	-2.21^{***}	-1.97^{**}		
	(0.67)	(0.90)		
6. Including other pay	-2.82^{***}	-2.61^{**}	-1.46^{***}	-0.34^{***}
	(0.81)	(1.10)	(0.34)	(0.12)

Notes.- second-step regression results of estimated period effects on unemployment rate, $\hat{\gamma}$. The first row is identical to Table 3, included here for comparison. The second row uses mean wages in jobs as the dependent variable in the first step. The third row excludes all time-varying controls from the first step. The forth row includes public sector firms in the analysis. The fifth row uses the Retail Price Index, instead of the Consumer Price Index, to deflate wages. The sixth row uses less restricted values of the first-step dependent variables: wages include shift-work, incentive payments, overtime, and all other payments; hours refer to basic and paid overtime.
Newey-West standard error estimates robust to first-order serial correlation in parentheses.
*** Statistically significant at the 1% level; ** at the 5% level, two-sided tests.

job stayers is larger. We prefer the median as a measure of the typical wage because it is less sensitive to changes in the extent of sampling error within jobs over time, given our specific sample selection criteria for jobs. The third row removes all controls for time-varying job characteristics from the first step. In doing so we would expect to underestimate the cyclical response of wages because of a procyclical composition bias along some observable characteristics. However the results here show that those observables that we do control for at the job level, namely gender, union coverage, age and firm size, are collectively not important in this regard. The fourth row includes jobs from the public sector. The main findings are qualitatively unchanged. The hiring hours in public sector entry-level jobs were somewhat more responsive to the Great Recession than in the private sector, potentially reflecting the squeeze on labour costs imposed by fiscal austerity. The fifth row simply illustrates the difference in results when we use an alternative price deflator. The RPI notably includes the cost of housing, including mortgage interest payments, whereas the CPI does not. Interest rates were cut during the Great Recession, and so the RPI is itself more cyclical than the CPI. Hence the measured real RPI-wage cyclicality is smaller, though still significant. We prefer the CPI because it is more internationally comparable and is the basis of the Bank of England's inflation target. The sixth row includes other work-related payments in earnings and the derived hourly wage rate, such as incentive or overtime pay. It similarly includes overtime and shift work in hours worked. There are reasonable arguments why including these other payments could lead to both increased or decreased wage responsiveness. Here their inclusion has no significant effect, except for job stayers' hours, suggesting that working hours for employees with overtime and shift work are more responsive than standard working hours.

TABLE B2: Estimated semi-elasticity of real wages and hours with respect to the unemployment rate, 1998-2016: sample selection robustness - varying the minimum number of employees per job-year required for inclusion in the CH-firms sample

	Wages		Hours	
Min. hires requirement	New hires (1)	Job stayers (2)	New hires (3)	Job stayers (4)
2 employees	-2.93^{***}	-2.22^{**}	-1.73^{***}	-0.03
	(0.74)	(0.90)	(0.54)	(0.17)
3 (baseline)	-2.83^{***}	-2.60^{**}	-1.47^{***}	-0.20
	(0.87)	(1.13)	(0.42)	(0.22)
4	-2.71^{***}	-2.44^{***}	-1.83^{***}	-0.16
	(0.76)	(0.91)	(0.56)	(0.13)
5	-2.72^{***}	-2.53^{***}	-1.43^{***}	-0.13
	(0.77)	(0.96)	(0.51)	(0.13)
6	-2.92^{***}	-2.18^{***}	-2.40^{***}	-0.36^{**}
	(0.68)	(0.81)	(0.67)	(0.18)
7	-2.94^{***}	-2.19^{**}	-2.23^{***}	-0.44^{**}
	(0.70)	(0.89)	(0.55)	(0.21)
8	-2.62^{***}	-2.24^{**}	-2.27^{***}	-0.52^{**}
	(0.66)	(0.94)	(0.65)	(0.22)
9	-2.54^{***}	-2.41^{**}	-2.75^{***}	-0.46
	(0.62)	(0.94)	(0.81)	(0.24)
10	-2.59^{***}	-1.96^{***}	-2.87^{***}	-0.48
	(0.56)	(0.74)	(0.92)	(0.27)

Notes.- second-step regression results of estimated period effects on unemployment rate, $\hat{\gamma}$. Each row gives results varying the minimum number of employees per job-year required for selection into the analysis sample. "3 (baseline)" is identical to Table 3, included here for comparison.
Newey-West standard error estimates robust to first-order serial correlation in parentheses.
*** Statistically significant at the 1% level; ** at the 5% level, two-sided tests.

TABLE B3: Estimated semi-elasticity of the difference in $\hat{\beta}_t$ series between new hires and job stayers with respect to the unemployment rate, 1998-2016

	Difference wages	Difference hours
Baseline sample	0.24	-1.38^{**}
	(0.32)	(0.36)
Table 3		
2. Including controls for share of full-time workers	−0.19	−0.57
	(0.31)	(0.32)
3. Job hires in at least 25% of years when the firm is observed	−0.16	−0.27
	(0.36)	(0.29)
4. All jobs observed in at least 2 years	0.41	−0.31
	(0.41)	(0.20)
5. Baseline sample, but weighted by number of employees per year	−0.27	-2.29^{**}
	(0.49)	(0.81)
Table 4		
2. Baseline with quadratic trend	−0.37	-1.29^{***}
	(0.17)	(0.37)
3. First differences (OLS)	0.20	0.31
	(0.52)	(1.09)
4. Baseline sample, but weighted by number of jobs per years	−0.29	-1.38^{***}
	(0.28)	(0.43)
Table 5		
1. Labour productivity (I) whole economy	-0.12^{**}	0.26^{***}
	(0.05)	(0.09)
2. Labour productivity (II) services sector	-0.14^{**}	0.28^{***}
	(0.05)	(0.09)

Notes.- second-step regression results of estimated period effects on unemployment rate, $\hat{\gamma}$. Dependent variable is the difference in composition-adjusted period means, $\hat{\beta}_t$, between new hires and job stayers.

Newey-West standard error estimates robust to first-order serial correlation in parentheses.

*** Statistically significant at the 1% level; ** at the 5% level, two-sided tests.

Appendix C. Additional tables

TABLE C1: Distribution of new hires over industries, all years 1998-2016

Industry (SIC2003)	Hires	Share
Wholesale and retail (52)	26,792	0.49
Accommodation and restaurants (55)	9,842	0.18
Financial intermediation (65)	3,821	0.06
Industrial cleaning and labour recruitment (74)	9,295	0.17
Other	5,467	0.10

Notes.- absolute and frequency distribution of new hires over industries. Shares might not sum to one due to rounding. Classification according to the ONS Standard Industrial Classification 2003.

TABLE C2: Distribution of new hires over occupations, all years 1998-2016

Occupation (ISCO88)	Hires	Share
Customer services clerks (41)	5,468	0.10
Personal and protective services workers (51)	6,561	0.12
Models, salespersons and demonstrators (52)	26,245	0.48
Sales and services elementary (91)	8,202	0.15
Labourers in mining, construction, manufacturing and transport (93)	2,734	0.05
Other	5,467	0.10

Notes.- absolute and frequency distribution of new hires over occupations. Shares might not sum to one due to rounding. Classification according to the ILO International Standard Classification of Occupations 1988.

TABLE C3: Estimated period-fixed effects for real hourly wages ($\hat{\beta}_t$ from first-step regressions)

	CH-firms		ASHE	
	New hires	Job stayers	New hires	Job stayers
Year	(1)	(2)	(3)	(4)
1998	−0.103	−0.084	−0.075	−0.093
1999	−0.087	−0.071	−0.083	−0.077
2000	−0.087	−0.051	−0.081	−0.052
2001	−0.046	−0.036	−0.044	−0.026
2002	−0.011	−0.005	−0.018	−0.007
2003	0.000	0.000	0.000	0.000
2004	0.006	−0.002	0.009	0.003
2005	0.022	0.015	0.024	0.018
2006	0.017	0.027	0.033	0.021
2007	0.031	0.023	0.033	0.019
2008	0.003	0.020	0.020	0.011
2009	0.023	0.017	0.020	−0.009
2010	−0.010	−0.022	−0.009	−0.050
2011	−0.054	−0.062	−0.033	−0.078
2012	−0.079	−0.099	−0.063	−0.114
2013	−0.087	−0.124	−0.074	−0.134
2014	−0.089	−0.136	−0.087	−0.144
2015	−0.069	−0.116	−0.069	−0.130
2016	−0.021	−0.088	−0.015	−0.112

Notes.- time series of period-fixed effects for different subsamples of the ASHE. Estimated using (1) and (3). Normalised to zero in 2003. (1) Entry-level new hires, (2) job stayers in CH-firms (3) ASHE new hires, (4) ASHE job stayers.

TABLE C4: Estimated period-fixed effects for basic weekly hours worked ($\hat{\beta}_t$ from first-step regressions)

Year	CH-firms New hires (1)	CH-firms Job stayers (2)	ASHE New hires (3)	ASHE Job stayers (4)
1998	0.063	0.004	0.014	−0.005
1999	−0.003	−0.001	−0.002	−0.003
2000	0.040	0.007	0.001	−0.001
2001	0.018	0.004	0.005	0.001
2002	0.021	0.002	−0.003	−0.001
2003	0.000	0.000	0.000	0.000
2004	0.003	−0.002	0.015	0.001
2005	−0.010	−0.010	−0.018	−0.011
2006	−0.027	−0.009	−0.015	−0.011
2007	0.016	0.005	−0.012	−0.013
2008	−0.052	0.010	−0.008	−0.011
2009	−0.053	0.003	−0.024	−0.017
2010	−0.059	0.003	−0.040	−0.015
2011	−0.131	−0.004	−0.049	−0.015
2012	−0.139	−0.009	−0.053	−0.016
2013	−0.154	−0.006	−0.048	−0.012
2014	−0.129	0.004	−0.046	−0.008
2015	−0.132	0.005	−0.051	−0.006
2016	−0.119	0.007	−0.062	−0.008

Notes.- see Table C3

TABLE C5: Time series of price deflators and business cycle indicators

Year	CPI	RPI	SPPI	Labour prod. whole economy	Labour prod. services sector
1998	95.56	92.48	97.58	90.62	91.73
1999	97.04	93.96	96.08	93.00	94.26
2000	97.58	96.78	95.97	95.99	96.37
2001	98.66	98.52	98.27	98.26	98.36
2002	100.00	100.00	100.00	100.00	100.00
2003	101.48	103.10	101.61	102.60	101.94
2004	102.69	105.68	103.34	105.67	104.52
2005	104.57	109.06	104.72	106.59	105.27
2006	106.72	111.89	108.06	109.46	108.60
2007	109.68	116.93	110.94	110.76	110.12
2008	112.90	121.84	115.09	112.34	111.79
2009	115.59	120.38	113.59	110.38	110.29
2010	119.89	126.76	115.67	109.80	109.88
2011	123.92	133.35	116.94	111.32	110.68
2012	123.03	138.02	117.86	112.15	111.48
2013	132.12	142.02	118.78	110.99	111.35
2014	134.54	145.57	120.51	110.65	110.87
2015	134.27	146.88	120.97	111.65	111.90
2016	134.68	148.79	122.81	112.42	112.62

Notes.- "CPI" - Consumer Price Index; "RPI" - Retail Price Index; "SPPI" - Services Producer Price Index; "Labour prod. whole economy" - chain volume measure of gross value added at basic prices in the UK; "Labour prod. services sector" - chain volume measure of gross value added at basic prices in services industries.

TABLE C6: Estimated hiring-year-tenure-fixed effects for real wages and hours relative to their hiring levels: workers who stay in entry-level jobs

		Wages			Hours		
Hiring year	Tenure:	1 year	2 years	3 years	1 year	2 years	3 years
2002		1.79	4.40	7.35	1.60	2.78	1.66
		(0.51)	(0.62)	(0.73)	(1.11)	(1.35)	(1.59)
2003		−0.09	2.63	6.68	0.80	2.62	2.09
		(0.60)	(0.73)	(0.87)	(1.30)	(1.59)	(1.88)
2004		2.07	7.13	8.61	0.91	2.79	0.95
		(0.54)	(0.67)	(0.81)	(1.17)	(1.46)	(1.76)
2005		4.86	6.89	8.78	1.37	0.46	3.54
		(0.48)	(0.61)	(0.74)	(1.04)	(1.33)	(1.60)
2006		2.26	2.67	4.53	−0.52	2.63	1.67
		(0.52)	(0.67)	(0.84)	(1.14)	(1.45)	(1.82)
2007		2.44	5.66	4.27	1.40	2.25	3.40
		(0.51)	(0.67)	(0.77)	(1.10)	(1.45)	(1.66)
2008		4.85	4.00	4.12	−1.20	2.08	2.13
		(0.49)	(0.59)	(0.68)	(1.06)	(1.27)	(1.48)
2009		−0.20	0.28	−0.80	1.25	3.55	4.72
		(0.47)	(0.59)	(0.79)	(1.02)	(1.28)	(1.72)
2010		1.11	−0.35	−0.37	0.98	1.06	3.20
		(0.54)	(0.74)	(0.90)	(1.17)	(1.61)	(1.96)
2011		−1.24	−1.65	−0.18	1.64	6.04	9.88
		(0.53)	(0.67)	(0.82)	(1.16)	(1.46)	(1.78)
2012		−0.39	1.17	5.79	4.63	8.28	14.64
		(0.51)	(0.67)	(0.86)	(1.11)	(1.45)	(1.86)
2013		3.04	7.70	14.30	8.75	16.35	18.46
		(0.54)	(0.72)	(0.94)	(1.17)	(1.55)	(2.04)

Notes.- see Section 5, Figure 6. Ordinary least squares standard errors in parentheses.

Appendix D. Additional figures

FIGURE D1: Characteristics of employees in the consistent-hiring-firms sample and whole ASHE: comparison of new hires in entry-level jobs vs. job stayers, 1998-2016

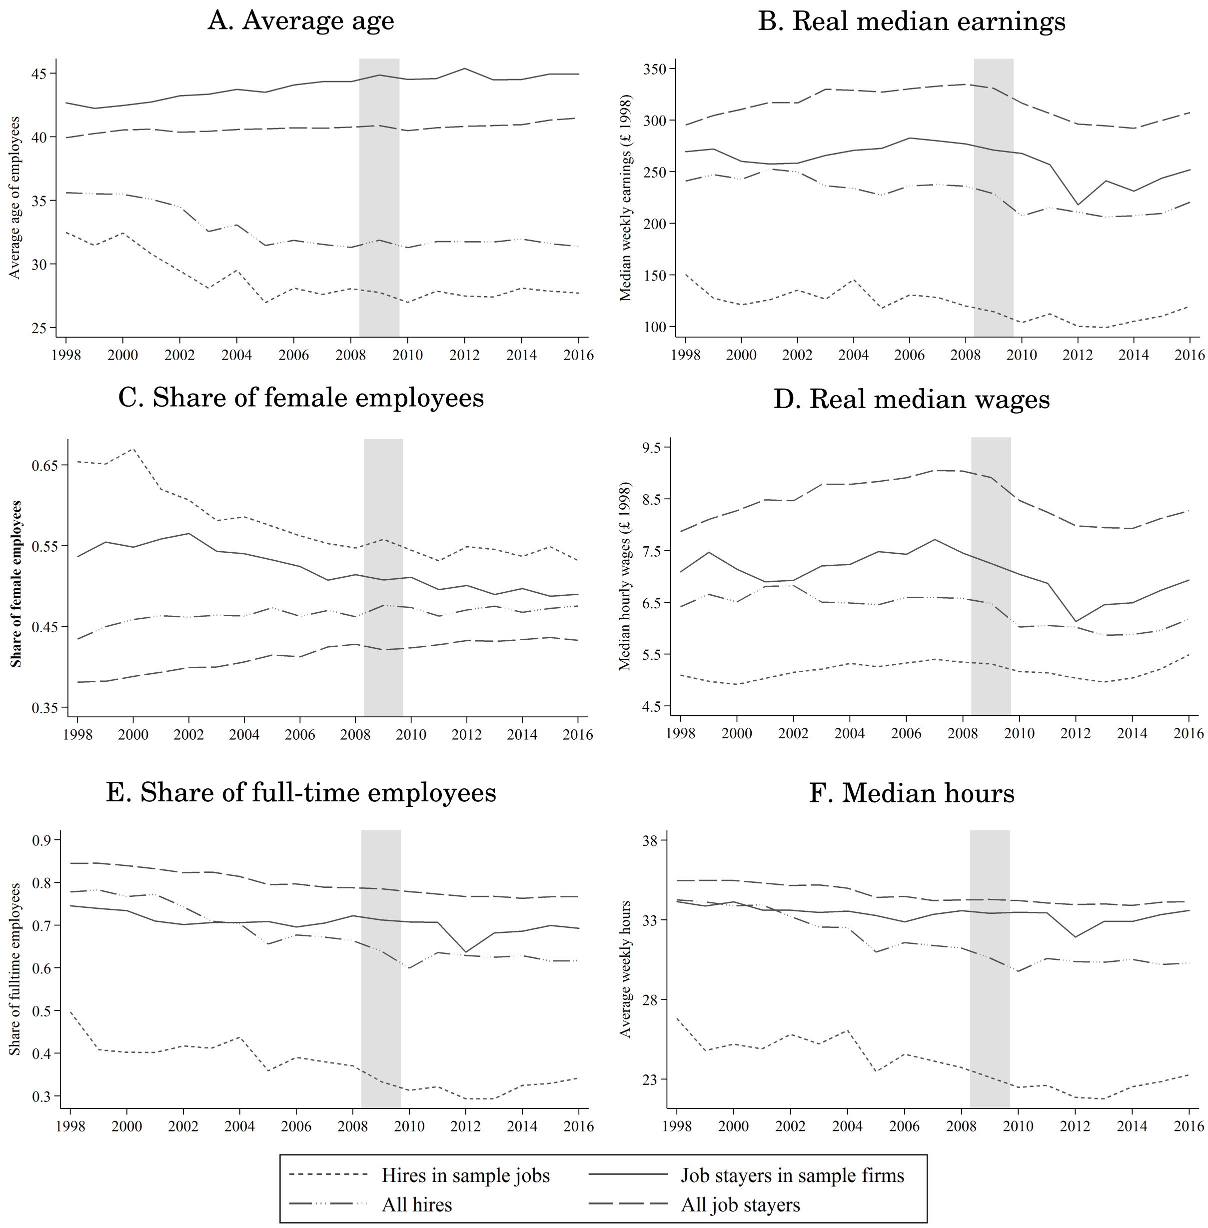

Notes.- shaded area marks official UK recession dates. “Hires in sample firms" refers to employees in entry-level jobs with less than twelve months tenure. “Job stayers in sample firms" are for jobs and employees who have more than 12 months tenure in the same job, and only for firms which are represented in the CH-firms sample. “All hires” and “All job stayers” show the corresponding series for new hires and job stayers in the ASHE, estimated as averages at the worker level. Ages 16-64 only.

FIGURE D2: Comparison of business cycle indicators, 1998-2016

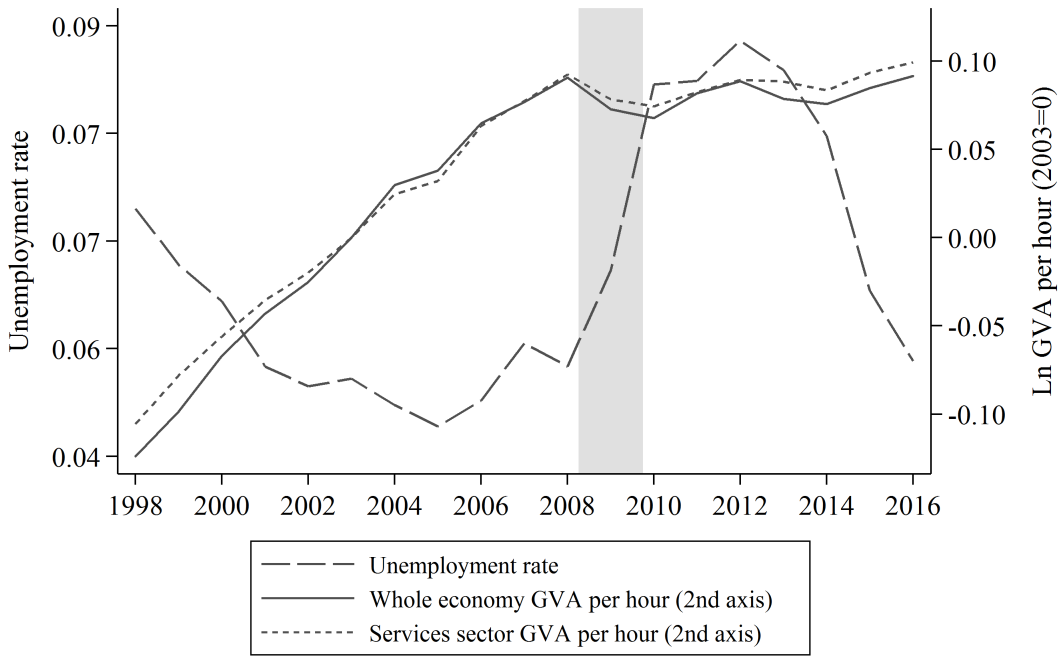

Notes.- see Section 2 and Table 5 for sources. Shaded area marks official UK recession dates.

FIGURE D3: Estimated period-fixed effects for real wages, 1998-2016: comparison of entry-level jobs, all new hires, job stayers in the CH-firms sample, and all job-stayers

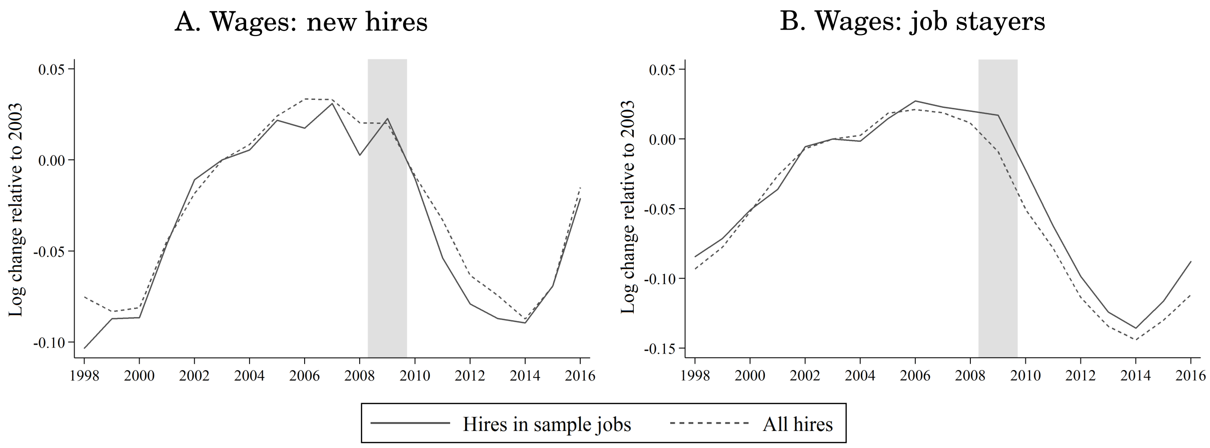

Notes.- see Figure D2 and Section 2 for further details of sample construction. “All” here refers to all firms and jobs represented in the ASHE. Shaded area marks official UK recession dates.

FIGURE D4: Estimated period-fixed effects for hours worked, 1998-2016: comparison of entry-level jobs, all new hires, job stayers in the CH-firms sample, and all job-stayers

A. Hours: new hires B. Hours: job stayers

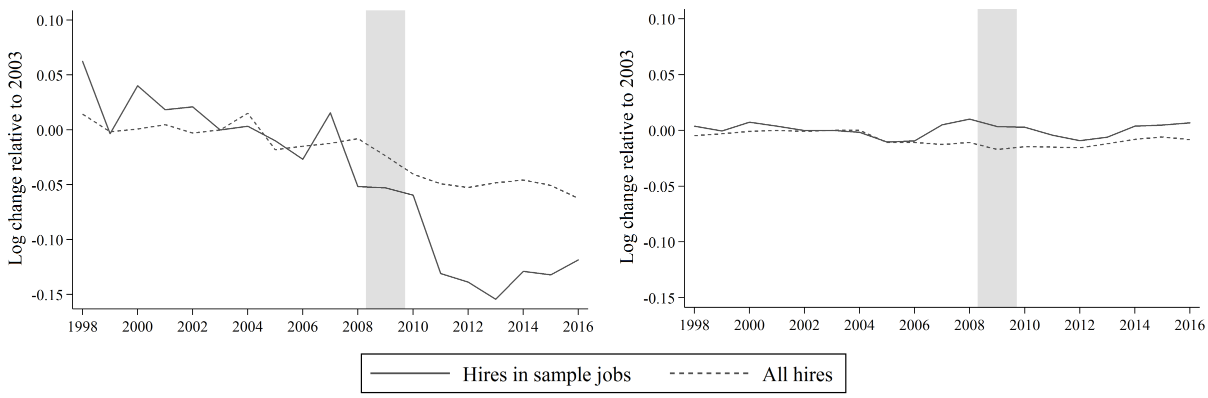

Notes.- see Figure D2 and Section 2 for further details of sample construction. “All” here refers to all firms and jobs represented in the ASHE. Shaded area marks official UK recession dates.

FIGURE D5: Estimated period-fixed effects for real wages including linear trend, 1998-2016: comparison of new hires in entry-level jobs vs. job stayers

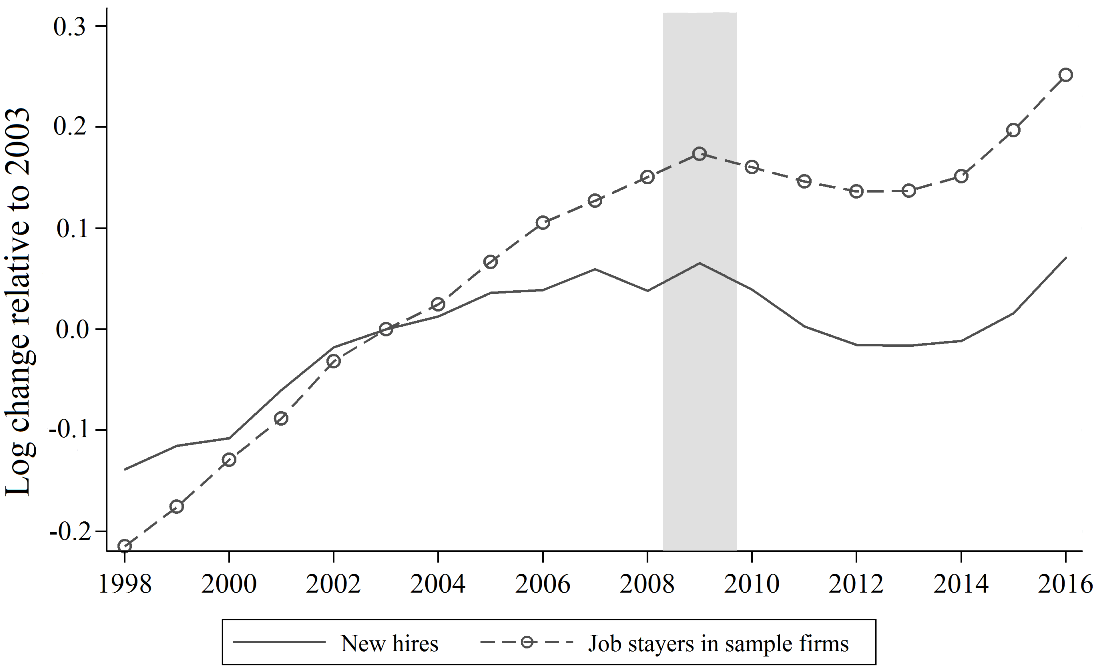

Notes.- see Figure 2, except here the series are not detrended. 2003 normalised to zero.

Appendix E. Nominal wage changes of job stayers in CH-firms

FIGURE E1: Distribution of year-to-year changes in log nominal hourly wages for job stayers in CH-firms, 2005-06 and 2009-10

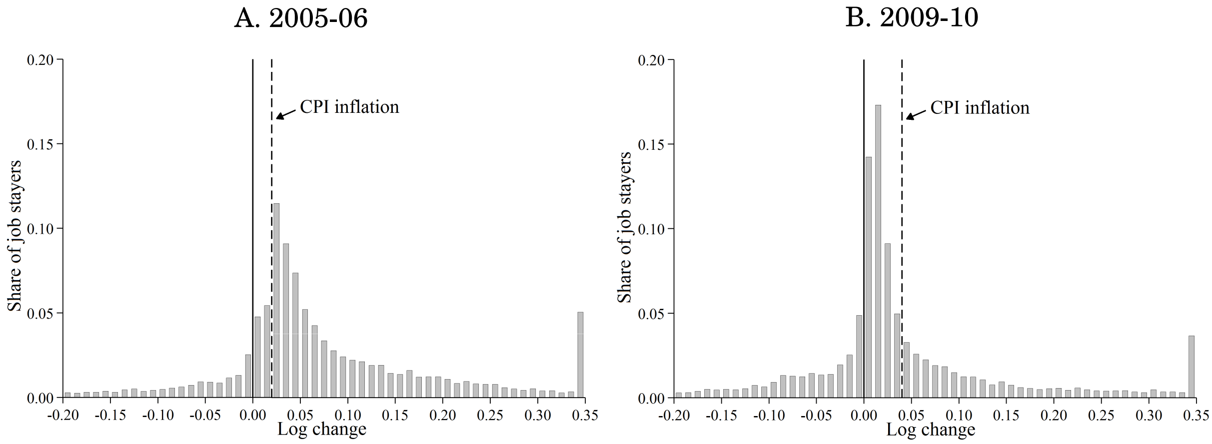

Notes.- solid line marks zero, dashed line marks the log change in the Consumer Price Index. Bars show half-open intervals, excluding the upper limit. Ages 16-64, private sector only.

TABLE E1: Percentage of job stayers with year-to-year change in log nominal hourly wages in given category, 1997-2016

	Percentage of log nominal wage change in given category					
Years	[−.01;0)	Exactly 0	(0; .01]	Nominal wage cut	Real wage cut	Inflation
1997-98	2.6	1.1	2.6	36.1	41.8	1.8
1998-99	1.7	0.9	2.2	22.8	27.2	1.5
1999-00	2.7	1.5	3.2	22.0	25.3	0.6
2000-01	2.5	1.4	2.8	18.0	22.5	1.1
2001-02	2.0	2.1	2.8	20.1	26.1	1.3
2002-03	2.9	0.7	3.5	23.7	29.2	1.5
2003-04	2.6	0.5	2.9	34.5	38.4	1.2
2004-05	2.6	0.8	3.3	20.0	27.5	1.8
2005-06	2.4	0.7	3.9	17.1	27.2	2.0
2006-07	2.4	1.7	3.2	21.7	37.2	2.7
2007-08	2.8	1.7	3.5	17.9	37.0	2.9
2008-09	3.4	2.9	4.2	20.1	34.3	2.4
2009-10	4.7	5.7	7.9	26.8	69.2	3.7
2010-11	3.8	3.7	5.7	24.5	61.2	3.3
2011-12	4.4	4.4	5.4	24.8	66.7	4.0
2012-13	3.8	4.3	5.3	24.7	55.8	2.4
2013-14	3.2	3.2	5.6	24.3	40.4	1.8
2014-15	2.7	1.5	4.7	18.9	18.0	-0.2
2015-16	3.2	2.4	4.3	25.9	29.6	0.3

Notes.- share of job stayers in CH-firms with log nominal wage changes in the indicated interval. Inflation is computed as average log change in CPI over previous four quarters.

Nominal wage changes of job stayers in the UK have been analysed previously by Nickell and Quintini (2003) and most recently by Elsby et al. (2016). We briefly summarise results for year-to-year changes in the log nominal wages of job stayers in our baseline sample of firms. Figure E1 shows the distributions of log changes for job stayers between 2005-06 and 2009-10. These two periods are representative of periods with relatively low (2005-06) and relatively high (2009-10) shares of job stayers with nominal wage cuts, see Table E1 which displays summary statistics for all years in our sample. The dashed line marks the inflation rate in the histograms. Bars in the histograms exclude upper limits, so log wage changes of exactly zero are included in the bin to the right of the solid line.

The spike at zero is relatively small during normal times, ranging from 0.5 per cent to 2.1 per cent in the period before the Great Recession as Table E1 shows. The distribution in panel A also suggests that most wages increase with the rate of inflation during normal times, thus keeping the real wage constant. Nevertheless, even during this period a notable share of job stayers, around 20 per cent, experienced nominal wage cuts. This share increased during the recession to around 25 per cent on average. Similarly, the share of job stayers with exactly zero nominal wage growth peaked at 5.7 per cent between 2009-10. In particular, Figure E1B displays a relatively large share of nominal wage changes between zero and two per cent for job stayers in CH-firms between 2009-10. However, the large increase in the percentage of job stayers which experienced negative changes in log *real* wages, as shown in the last column of Table E1, was mainly caused by the rise in inflation.

These findings suggest that zero is a significant threshold for nominal wage changes and limited the downward adjustment of nominal wages, as Nickell and Quintini (2003) argue, but on average more than 20 per cent of job stayers experience nominal wage cuts in the UK, suggesting that there is a relatively high degree of nominal wage flexibility in the British labour market. Nevertheless, the increase in inflation during the Great Recession resulted in over two-thirds of job stayers seeing their real wages cut.

Appendix F. Description of the kernel re-weighting method

This section describes the method of DiNardo et al. (1996), which we use to estimate the counterfactual densities of real hiring wages: the exposition here follows closely theirs.

Let $f^i(w|x;m_i)$ be the density of real hiring wages in period i, conditional on observable attributes x and the real minimum wage m_i. The density of observed attributes in period i is $h(x|t=i)$. The observed densities of real hiring wages in two periods, say 2004 ($i=04$) and 2013 ($i=13$), are

$$g(w|t=04;m_{04}) = \int_{\Omega_x} f^{04}(w|x;m_{04})h(x|t=04)dx \ , \tag{5}$$

and

$$g(w|t=13;m_{13}) = \int_{\Omega_x} f^{13}(w|x;m_{13})h(x|t=13)dx \ , \tag{6}$$

where Ω_x is the domain of observed attributes. Differences in attributes at the job-level between the two periods are captured by the density functions $h(x|t=04)$ and $h(x|t=13)$. Differences in the "price" paid for these attributes are captured by differences in $f^{04}(w|x;m_{04})$ and $f^{13}(w|x;m_{13})$, and these differences can depend on the real minimum wage.

The counterfactual density of real hiring wages that would prevail if the level of the real minimum wage of 2004 was realised in 2013, *and* prices had remained at their 2013 level, is

$$g(w|t=13;m_{04}) = \int_{\Omega_x} f^{13}(w|x;m_{04})h(x|t=13)dx \ , \tag{7}$$

where we know $h(x|t=13)$, but the density of prices $f^{13}(w|x;m_{04})$, consisting of the real wage schedule of 2013 and the real minimum wage of 2004, is unobserved. We can partition the density in (7) into the part of real wages below the real value of the minimum wage in 2013 and the part of real wages above this threshold:

$$\begin{aligned} g(w|t=13;m_{04}) = &\int_{\Omega_x} [1-I(w \le m_{13})] f^{13}(w|x;m_{04})h(x|t=13)dx \\ &+ \int_{\Omega_x} I(w \le m_{13}) f^{13}(w|x;m_{04})h(x|t=13)dx \ , \end{aligned} \tag{8}$$

where $I(w \le m_{13})$ is an indicator function that equals one if the observed wage is at or below the level of the real minimum wage in 2013. We follow DiNardo et al. and make quite restrictive economic assumptions, but because of this restrictiveness, they are also transparent.

Assumption 1 *Between two periods $i=\{L,H\}$ with $m_L < m_H$, the conditional density of hiring wages above the real value of the minimum wage m_H is not affected by the minimum wage:*

$$[1-I(w \le m_H)] f^i(w|x;m_H) = [1-I(w \le m_H)] f^i(w|x;m_L) \ . \tag{9}$$

This is a conservative assumption, and we later conduct a sensitivity analysis where we allow the minimum wage to affect the wage density above its real value. Our results vary with the size of this spillover effect, but not substantially.

Assumption 2 *The shape of the conditional density at or below the minimum wage depends only on the real value of the minimum wage. Thus, the conditional density in $i = L$ below m_H is proportional to the conditional density in $i = H$ below m_H:*

$$I(w \leq m_H) f^H(w|x; m_L) = I(w \leq m_H) \psi_w f^L(w|x; m_L)\ , \tag{10}$$

where the re-weighting function ψ_w will be defined below.

With this assumption we can regard the density in 2004 below the real value of the 2013 minimum wage as the latent hiring wage distribution, conditional on observable attributes x. Without a structural model, it is not possible to impute the wage schedule below the real minimum wage in 2013 without making strong assumptions like Assumption 2, but we think this assumption is at least relatively transparent. The last assumption necessary to derive a counterfactual hiring wage density is:

Assumption 3 *The level of the minimum wage can affect the number of new hires, but has no effect on the number of entry-level jobs.*

This assumption is weaker than the corresponding one originally made by DiNardo et al. (1996), who assumed that the level of the minimum wage does not affect the level of employment.

We illustrate the method in Figure F1, which shows estimated densities of real hiring wages (we pool data over multiple periods in Figures F1 and F2 for data confidentiality reasons). The kernel density estimate in the top panel uses pooled data from 2002-04, a period where the NMW was relatively less binding. The solid vertical line shows the real value of the NMW at its 2013 level. Panel B shows the corresponding density for 2012-14, where the real NMW remained nearly constant around its 2013 value. The counterfactual density displayed in panel C is a simple combination of the part of the density to the left of the solid line in panel A and to the right of the solid line in Panel B, scaled to integrate to one.

Increasing the assumed spillover of the NMW acts as if shifting the NMW in panels A-C to the right: the area of the density below the new threshold, consisting of the NMW plus spillover, will increase. This means that a larger part of the 2012-14 density will be replaced by the 2002-04 density. In the extreme case that the NMW plus spillover exceeds the highest measured job-level hiring wage in 2012-14, the counterfactual density would fully consist of the estimated density in 2002-04.

FIGURE F1: Illustration of the re-weighting procedure for log real hourly hiring wages

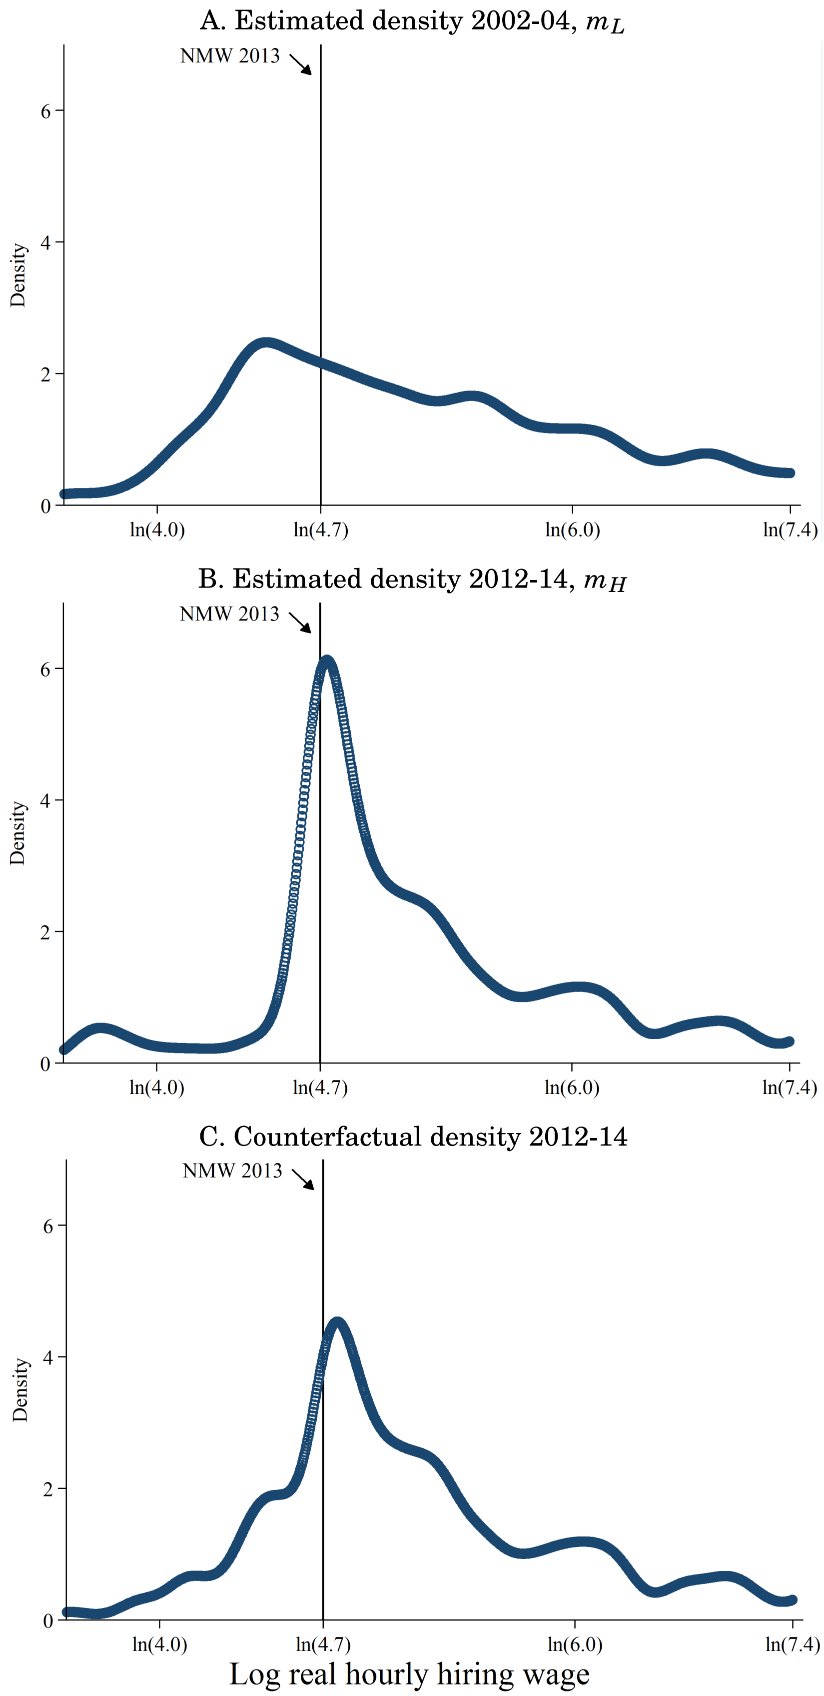

Notes.- densities estimated using Gaussian kernel and bandwidth of 0.03 (A) and 0.02 (B-C). Monetary values deflated to 1998 values using the CPI. Solid lines show the real value of the adult rate minimum wage in 2013.

Figure F2D plots the estimated (connected circles) and counterfactual (solid line) density for 2012-14 together. Most of this mass originates from jobs which are observed to hire slightly above the NMW, a result of the smoothing by the kernel estimator. For hiring wages which exceed the NMW substantially, the estimated and counterfactual density are, as expected, indistinguishable. Figure F2E displays the difference between the estimated and counterfactual wage densities shown in Panel D of this figure. The difference is negative for values around the value of the NMW in 2013 and positive for log hiring wages between ln(4) and ln(4.6).

FIGURE F2: Estimated and counterfactual densities of real log hiring wages at the job-level, 2012-14

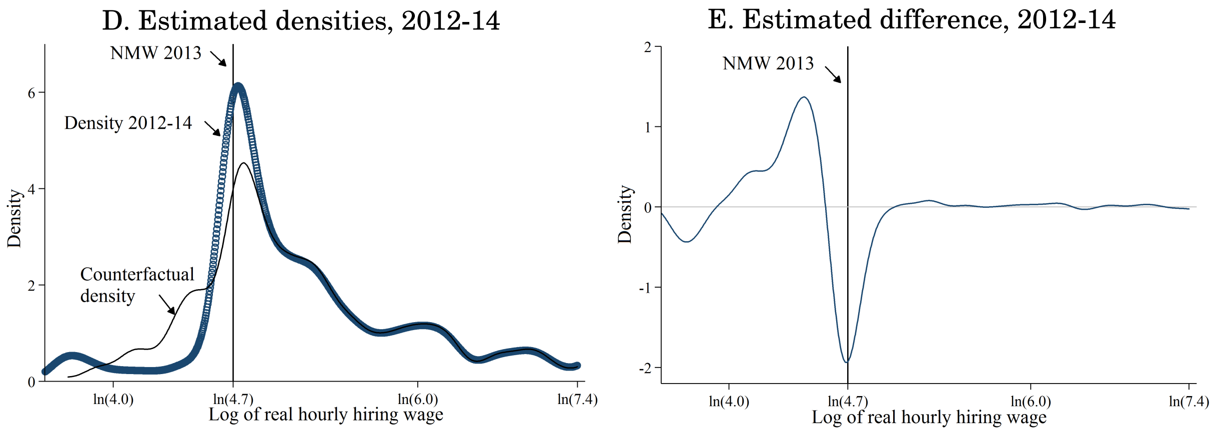

Notes.- densities estimated using Gaussian kernel and bandwidth of 0.02. Monetary values deflated to 1998 values using the CPI. Solid vertical line shows the real value of the adult rate minimum wage.

Assumptions 1-3 allow us to write

$$\begin{aligned} g(w|t=13;m_{04}) = & [1 - I(w \le m_{13})] f^{13}(w|x;m_{04}) dx \\ & I(w \le m_{13}) \psi_w f^{04}(w|x;m_{04}) dx \;, \end{aligned} \tag{11}$$

with

$$\psi_w = \frac{\Pr(w \le m_{13}|x, f = f^{13})}{\Pr(w \le m_{13}|x, f = f^{04})} \;, \tag{12}$$

which ensures that the density integrates to one over the distribution of attributes. The counterfactual real hiring wage density is found by integrating over the observed distribution of attributes:

$$\begin{aligned} g(w|t=13;m_{04}) = & \int_{\Omega_x} [1 - I(w \le m_{13})] f^{13}(w|x;m_{04}) h(x|t=13) dx \\ & + \int_{\Omega_x} I(w \le m_{13}) \psi_w f^{04}(w|x;m_{04}) h(x|t=13) dx \;. \end{aligned} \tag{13}$$

The key insight of DiNardo et al. is that the wage density that would result from combining the wage schedule in 2004, $f^{04}(w|x;m_{04})$, and the marginal distribution of attributes, $h(x|t=13)$, can be obtained by taking the observed density of attributes in 2004, $h(x|t=04)$, and re-weighting it to reflect differences between the two periods. Let this re-weighting function be denoted θ, then

$$g(w|t=13;m_{04}) = \int_{\Omega_x} [1-I(w \le m_{13})] f^{13}(w|x;m_{04}) h(x|t=13)dx \\ + \int_{\Omega_x} I(w \le m_{13}) \psi_w f^{04}(w|x;m_{04}) \theta h(x|t=04)dx \ , \tag{14}$$

where the re-weighting function is

$$\theta = \frac{h(x|t=13)}{h(x|t=13)} = \frac{\Pr(t=13|x)}{\Pr(t=04|x)} \frac{\Pr(t=04)}{\Pr(t=13)} \ . \tag{15}$$

The last equality follows from Bayes' rule. We combine the two re-weighting functions to

$$\psi = \theta \cdot \psi_w = \frac{\Pr(t=13|x, w \le m_{13})}{\Pr(t=04|x, w \le m_{13})} \frac{\Pr(t=04)}{\Pr(t=13)} \ . \tag{16}$$

We estimate the probability of a job being below the NMW in 2013, conditional on its observed attributes parametrically, using a logit model

$$\Pr(t=13|x, w \le m_{13}) = \Lambda(C(x)) \ , \tag{17}$$

with $C(x)$ being a vector that is a function of the covariates x. The covariates are: one-digit industry dummies, a cubic in age and firm size, and the shares of workers that are female, full-time, permanent, and covered by a collectively bargaining agreement. We then compute estimates of $\hat{\psi}$ for each observation, and use these weights in the kernel density estimation to derive the counterfactual density of real hiring wages in 2013. In particular, the weight equals one if an observation is above the NMW in 2013, it equals zero if the observation is below the NMW in 2013, and the weight equals $\hat{\psi}_j$ if an observation j in the pooled data is from 2004 and from the section of the density of real wages below the real value of the NMW in 2013.