Cyclical labor costs within jobs

Cyclical Labor Costs Within Jobs[*]

Daniel Schaefer[†]
University of Edinburgh

Carl Singleton[‡]
University of Reading

August 2019

Abstract

Using UK employer-employee panel data, we present novel facts on how wages and working hours respond to the business cycle within jobs. Firms reacted 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 2.6 percent for new hires as well as for job stayers. Hiring hours worked were substantially procyclical, while job-stayer hours were acyclical. These results show that real wages are not rigid and that the labor costs of new hires are especially flexible.

Keywords: Wage rigidity; Great Recession; Hours worked; Job-level analysis
JEL codes: E24, E32, J30

[*]We are grateful for helpful comments and advice from two anonymous referees, Steven Dieterle, Mike Elsby, Marco Fongoni, Jan Grobovšek, Andy Snell, Andreas Steinhauer, Jonathan Thomas, Ludo Visschers, and participants in 2017 at the Conference on Markets with Search Frictions (Aarhus), 8th ifo Conference on Macroeconomics and Survey Data (Munich) and 42nd Simposio de la Asociación Española de Economía (Barcelona), and in 2018 at the 8th Search and Matching Annual Conference (Cambridge) and Royal Economic Society Annual Conference. 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. A previous version of this work was published as CESifo Working Paper No. 6766 in November 2017 and titled "Real wages and hours in the Great Recession: Evidence from firms and their entry-level jobs".

[†]School of Economics, University of Edinburgh, 31 Buccleuch Place, Edinburgh, EH8 9JT, UK; daniel.schaefer@ed.ac.uk

[‡]Corresponding author: Department of Economics, University of Reading, Edith Morley Building, Reading, RG6 6UB, UK; c.a.singleton@reading.ac.uk; with thanks to the Economic and Social Research Council (UK) for funding support under Grant No. ES/J500136/1

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 labor 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 responses of wages and hours *within* jobs, which is what determines a firm's employment decisions in frictional labor markets. For instance, suppose that workers switch from high- to low-paying jobs during recessions. Then aggregate wages would decline, even if the wages which firms pay their employees within jobs are completely rigid. Unlike previous studies for the UK, we use a longitudinal employer-employee dataset, which allows us to measure the business cycle response of real hourly wages and weekly hours worked *within* jobs between 1998 and 2016.

Our main contribution to the wider literature is to combine the robust measurement of job-level 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, UK firms significantly reduced the real hourly wages of new hires and job stayers within jobs during the most recent 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 is associated with average decreases of 2.6 percent in the real hourly wages of both new hires and job stayers. The weekly hours worked of new hires respond to the same increase in unemployment by an average decrease of 1.7 percent. A shift from full- to part-time work explains the majority of the cyclical decline in the hours of new hires. However, there is no significant difference between the wage responses of full- and part-time workers within jobs.

In a wide class of labor market models, firms' employment decisions are forward-looking and depend on expected labor costs (Kudlyak, 2014). Therefore, we track new hires for up to three years in continuing job matches to understand how persistent their initial hiring wages and hours are. We find strong cohort effects: accounting for unobserved match quality, real wage growth on the job decreased to zero for the cohorts of employees hired during the Great Recession, and this stagnant wage growth was only partially compensated by larger annual increases in working hours for employees who stayed in the same job for up to three years. These findings suggest that the sum of real wage payments in a job-match over time, i.e. the present value of labor costs for new hires, is even more responsive to business cycle conditions than the initial hiring conditions.

We use a simple empirical approach to measure responses to the business cycle, proposed by Martins et al. (2012). To obtain "typical" job-level measures, 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 with job-fixed effects. We identify a sample of jobs into which firms consistently hired before, during and after the recession. This matters because we use within-job variation to measure responses to the business cycle. If jobs with relatively rigid wages and hours simply stopped hiring during the most recent downturn, then the sample would over-represent jobs with particularly flexible hiring conditions. This approach trades off representativeness of the whole economy for confidence that economically meaningful responses are being estimated, at least from the perspective of what matters to firms. Looking at the job-level implies that the number of employees per job over time does not affect the estimated responses of wages and hours to the business cycle. This leaves results not only robust to cyclical changes in the composition of hiring but also any job-related secular trends in the labor market, such as the occupational polarization of employment in the UK (e.g. Goos and Manning, 2007; Salvatori, 2018). Using this approach, Martins et al. (2012) find that the real wages of new hires in Portugal decrease significantly by 1.8 percent when the unemployment rate increases by one percentage point, lower than our similarly obtained UK estimate of 2.6 percent.[1]

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

Recent work finds that UK real wages are procyclical for job stayers, with an especially large response to the Great Recession (Gregg et al., 2014; Elsby et al., 2016). There exists some evidence that the wages of workers who change employers respond to the business cycle (see Bils, 1985; 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. However, those findings could be masking wage rigidity within jobs if workers transition from high-paying jobs to low-paying jobs during recessions, and vice versa during expansions.

[1] Also accounting for cyclical job switching, Carneiro et al. (2012) and Stüber (2017) find significant wage flexibility in Portugal and Germany, respectively, though from estimating this at the level of workers rather than jobs. Stüber also looks at the job-level approach, estimating that the real daily wages of new hires and job stayers decline by 0.9 percent if the unemployment rate increases by one percentage point.

The sample of jobs we select and study mostly has high employee turnover and low wages, and is representative of around 50 percent of UK employment. Since the employment of low-wage workers typically declines sharply during UK recessions (Blundell et al., 2014), understanding why this group's employment did not drop more following the 2007/8 global financial crisis is of particular relevance.

All the aforementioned related studies focus on real wages. But firms can also adjust per worker labor costs by decreasing employee hours worked. The significant decline in UK 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, 2019), though not at the job level. The data used here allows us to examine the response of weekly hours worked within jobs, which previous studies could not address. A novel finding is that firms responded to the Great Recession by significantly decreasing the hours of new hires; whereas the hours worked by job stayers were not responsive. To the best of our knowledge, the relatively greater and large response of hiring hours to business cycle conditions has not been documented robustly before.[2] This new evidence of firm behavior, and its wider economic implications, should be further studied outside the specific context of the UK's experience of the Great Recession.

2 Methodology and data: measuring wage cyclicality within jobs

To measure the response of real hourly wages and weekly hours worked to the unemployment rate, we use a two-step regression approach (Solon et al., 1994; Martins et al., 2012). Compared with the alternative one-step approach, the results are more transparent and do not rely on asymptotic theory to correct for correlation across jobs within periods. We expand on this 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 hourly wage of new hires in some 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 some time-varying job-level characteristics in the vector $\boldsymbol{x}_{jt}$. We include these covariates to control to some extent for possible changes

[2] Using a British household panel dataset, Blundell et al. (2008) find that employees typically adjust their working hours in response to welfare reforms by switching between firms, and this is particularly the case for larger firms and in the services sector. Their data don't allow them to study the responsiveness of hours worked within jobs.

in the composition within jobs of employees and their characteristics over the business cycle, which we discuss further later.

The parameter estimates $\widehat{\beta}_t$ obtained from Equation (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 business cycle by regressing it on the unemployment rate U_t:

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

We measure the semi-elasticity of real wages with respect to the unemployment rate (or some other cyclical indicator) by the coefficient estimate $\widehat{\gamma}$. 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.[3] 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 robust to cross-sectionally correlated errors (or "cluster robust") with relatively few periods: the only requirement is that the sample mean $\widehat{\beta}_t$ is a good approximation of the population mean β_t. We later vary the exact specification of both steps for robustness, but the baseline second step also allows for a constant and a linear time trend.

To obtain a comparable series of $\widehat{\beta}_t$ for job stayers, we alter the first step. Let w_{jkt} be the median real hourly wage in some job j for stayers in some consecutive time period k. In our analysis, stayers are always defined as employees who work in the same job for two consecutive years, t and $t-1$, i.e. job stayers who work in the same occupation-firm pair j but in different pairs of years, e.g. 1998-9 and 2008-9, have different values of k. We use least squares to estimate:

$$w_{jkt} = \alpha_{jk} + \beta_t + \boldsymbol{x}'_{jkt}\boldsymbol{\delta} + \varepsilon_{jkt} \ , \tag{3}$$

where α_{jk} are job-two-year-period fixed effects and $\boldsymbol{x}'_{jkt}$ contains time-varying job-level characteristics as above.[4] The second-step regressions for job-stayer and hiring wages

[3] To illustrate, 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 within a period.

[4] Note that covariates such as the average age or tenure of the employees in a job, which could be included in the new hires first step, must be omitted here as they are collinear with the period-fixed effects, as they increase between periods by one for all stayers. To address this, we estimate and remove a linear trend from the estimated time series of $\widehat{\beta}_t$ for job stayers, before comparing it graphically with the equivalent series for new hires and normalizing both (see Figure 1).

are always identical, using the derived comparable estimates of $\widehat{\beta}_t$ in both cases as the dependent variable.

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

2.1 The Annual Survey of Hours and Earnings and other data

The Annual Survey of Hours and Earnings (ASHE), 1997-2016, is based on a one percent random sample of employees, drawn from the UK tax collection authority's (HM Revenue and Customs) Pay As You Earn (PAYE) records, which is the UK income tax withholding system. Questionnaires are sent to employers, who are legally required to complete them with reference to payrolls for a certain week in April. The ASHE is generally considered to provide accurate records of pay components (Nickell and Quintini, 2003).

The dataset is a panel of employees without attrition, forming an approximate one percent random sample of UK employees in every year.[5] Particularly valuable for our analysis are the longitudinal identifiers for individuals (1997-2016) and enterprises (2003-16). 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 organizations pay-setting practices are determined at the enterprise level.[6] For further information on how we construct an employer-employee panel from ASHE cross-sections, the employee, firm and job-level variables used, other adjustments made to the data and the sample selection, see Online Appendix A.

The analysis focuses on the two main components of employee remuneration: basic weekly paid hours and the hourly wage rate, which latter variable 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).[7] 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 (e.g. sick leave).

[5]The two main reasons why individuals might not be observed in some year are that they are non-employed or have changed employer between January and April. Since the questionnaires are in most cases sent in April to the employer's registered address according to January PAYE records, workers who switch employers during these months are undersampled.

[6]Brown et al. (2003) find that pay-setting in large UK companies mostly takes place at the enterprise level: in half of these companies, corporate management was determining pay directly, while in one-third corporate management was establishing the limits within which local managers had to negotiate.

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

The main indicator variable used for the business cycle is the working-age unemployment rate: the number of people unemployed divided by the economically active population.[8] 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 used. 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 focus on the unemployment rate for comparability with the wider literature. In Section 4 we discuss the robustness of the main results to this choice.

2.2 Constructing a 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 the sample to observations for the years 2003-16, because for this period we have almost complete records of firm identities and employment start dates. We exclude all firms which are observed for less than three years, as well as any in the public sector. 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.[9] The main results do not change substantially when we require tenure of less than six months, but the sample size is significantly smaller.

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 percent random employee sample, these requirements impose an effective lower bound on firm size in the entry-level job samples. Of the firms in the baseline sample, 95 percent 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 identities before 2003. In what follows, we refer to this sample of firms as consistent-hiring-firms (CH-firms). The selection criteria are naturally somewhat arbitrary, though hopefully reasonable. We will vary them for robustness when discussing the main empirical results, to reassure that the way the sample of jobs is selected does not drive the main results. For instance, it is certainly a limitation that the analysis will focus mainly on jobs within large UK firms. However,

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

[9] The ASHE questionnaire asks the following question: "When did this employee start working for your organisation [month, year]? If the employee has worked in another part of the organisation, or the organisation has changed ownership since the employee first joined, the start date should be the date when they first started work in the organisation. If this employee has left and was then re-employed, the start date should be the date they were re-employed."

we will also present some results representative of the whole UK labor market, with the caveat that we cannot then address the selective mobility of workers between jobs over time.

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

The baseline entry-level jobs sample consists of 347 firms hiring into 391 jobs (Table 1). The 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 can imply endogenous sample selection. The contemporaneous correlation of the number of entry-level jobs in the sample and the unemployment rate is insignificant (p-value: 0.49), and no other cyclical patterns are evident in Table 1 column (2). The median number of new hires per entry-level job and year is seven over the sample period.

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	4,176	
Unique	48,744	391	347	

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

A contribution of this paper is that we analyze the real wages and hours worked 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, i.e. who have job-specific tenure of at least 12 months, and 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. This may include stayers in both entry-level jobs and other jobs within the same firms. The sample consists of 7,779 repeated observations of occupation-firm pairs, totaling 158,194 job stayers. The selected jobs represent approximately 90 percent of all job stayers in the CH-firms sample over the whole period.

New hires in the baseline 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 basic hours of new hires are lower than for job stayers. The same statements hold for the entire ASHE, which represents the whole non-public sector economy, (columns (3) and (4), Table 2), though the difference between the hours worked by all new hires and job stayers is therein considerably smaller: almost two-thirds of hires into entry-level jobs are part-time. The lower average age, real wages, and basic hours in the baseline sample can be explained by the differences in industry and occupation compositions. Over two-thirds of new hires are made by firms in the "accommodation and restaurant" and the "industrial cleaning and labour recruitment" industries. Similarly, the largest shares are employed as service or sales workers (see Online Appendix Tables C1-C2 for complete industry and occupation breakdowns). This is also reflected in so far as large firms dominate the baseline sample. These large firms have average annual growth in the number of their employees of around four percent, while the average for all firms is around eight percent. The job-level regression approach we use here assumes that job and worker characteristics within jobs are invariant to business cycle fluctuations. We can partially address this by including controls for these characteristics in Equation (1). We are also reassured that although the observable characteristics of new hires in entry-level jobs and job-stayers within the same firms exhibit secular trends during 1998-2016, we do not see any notable cyclical patterns (Online Appendix Figure D1).[10]

Since our subsequent analysis is at the job level, we compute median wages and hours within jobs each year. Although we can expect some dispersion around these median wages, the robustness of the econometric approach and the meaningful interpretation of any results to some extent depend on us capturing "typical" hiring wages, given the one percent sample of employees. More than 50 percent of hiring wages lie within a range of five log points, and almost 90 percent within 10 log points, of their associated job-specific median values (Online Appendix Figure D2). Moreover, we did not find evidence of systematic variation in the distribution around the "typical" hiring wages with the business cycle. In the robustness discussion below, we show that the main

[10]Online Appendix Figure D1 shows that the share of men among new hires increases steadily by around ten percent from 1998 to 2016, while the share of full-time employees decreases. Because of the 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 analysis does not change the results.

TABLE 2: Descriptive statistics for employees and firms: 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)	7,767	7,767	45	29
Firm size annual growth	4.3%	4.3%	7.9%	7.9%
Total 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. Firm size annual growth is measured as an unweighted average across firms in the sample, rather than over employees.

results do not notably change if we use mean wages or hours within jobs instead. We also consider the effects of using values at the 25th or 75th percentiles within jobs.

For comparison with the wider literature and the reader's general interest, we include statistics and discussion concerning nominal wage changes and rigidity within the baseline sample in Online Appendix E.

3 Regression results

3.1 Job-level responses to the UK unemployment rate

Figure 1 shows the estimated time series $\widehat{\beta}_t$ for new hires and job stayers in CH-firms from regression models (1) and (3) over the period 1998-2016, after subtracting linear trends for comparability. The short-dashed lines indicate 95 percent confidence intervals for the point estimates of $\widehat{\beta}_t$ for new hires, using standard errors robust to serial correlation at the job-level. All series are normalized to zero in 2003. We estimate Equation (1) using the unbalanced baseline panel of jobs described in Table 1, where the baseline set of job-level time-varying covariates are: cubic functions in age and firm size, the share of female employees and the share of employees covered by a collective agreement.[11] These time series should be interpreted as composition-adjusted real wage rates and hours worked of new hires and job stayers at the job-level.

Figure 1A shows that real hiring wages increased by 13 log points between 1999 and 2007, while job-stayer wages increased by almost 14 log points among the same firms,

[11]The results remain virtually unchanged when we instead use dummies for ranges of age and firm size, instead of the cubic, implying that the model is well-specified while preserving degrees of freedom.

FIGURE 1: Estimated period-fixed effects for log real wages and log hours worked, including 95% confidence intervals for new hires, 1998-2016.

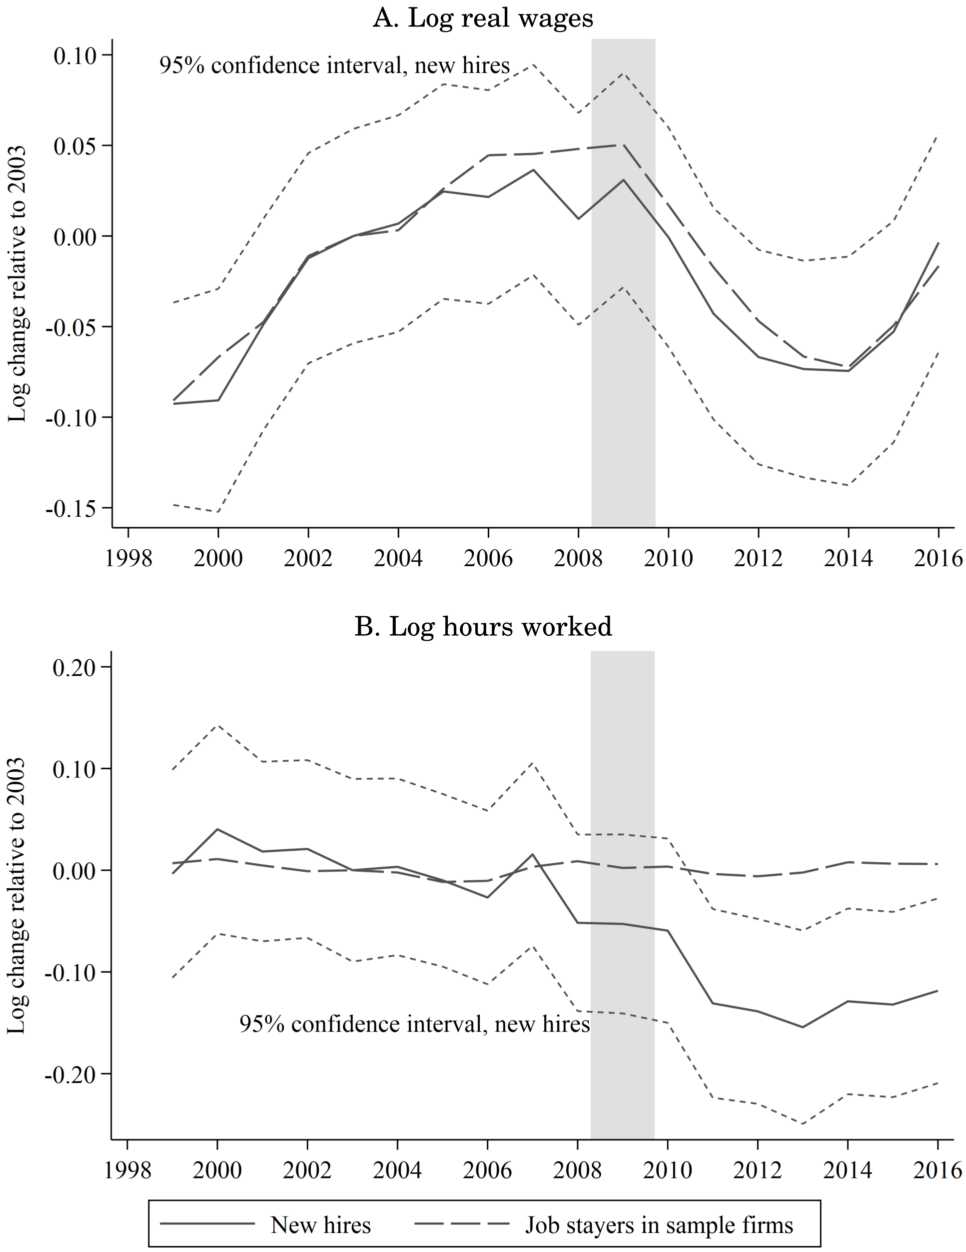

Notes.- the 95% confidence intervals for job stayers (not shown) are narrow and almost indistinguishable from the long-dashed lines. Standard errors are robust to clustering at the firm-level. Excluded reference category in first-step regressions (1)-(3) is 1998. Series normalized to zero in 2003. For comparability, linear trends have been removed. 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" 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.

relative to their trend. During the Great Recession, hiring wages then decreased by 11 log points until 2014, before slightly recovering over the next two years. The real wages of job stayers plummeted by around 12 log points between 2007 and 2014. For comparison, Elsby et al. (2016) document a decline in job-stayer real hourly wages in the whole UK economy between 2008 and 2012 of 14 log points for men and 8 for women, based on worker-level estimates. In Online Appendix F, we present evidence that the comparatively smaller drop of new hires' wages during the recent recession could be explained by the UK's National Minimum Wage (NMW). Intuitively, because job stayers' wages are typically higher than new hires' wages, the latter had less space to fall before reaching the NMW. Figure 1B shows the estimated series for hours worked among the same employees, jobs and time period. Hiring hours decreased by 15 log points between 2007 and 2014, being approximately constant before and after that period. In contrast, the hours worked by job stayers saw no significant change during the Great Recession.[12]

We measure the response of real wages and hours to the unemployment rate 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 the baseline regressions, because the ordinary least squares (OLS) residuals do not display significant evidence of heteroskedasticity.[13] 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 2.6 percent if the unemployment rate increases by one p.p.[14] These estimates are both significantly different from no response and they do not significantly differ from one another (see Online Appendix Table B3 for evidence on the significance of the differences between the estimated hiring and job stayer responses discussed throughout the paper). In column (3), we find that hiring hours have a semi-elasticity estimate with respect to the unemployment rate of around -1.7, compared to only -0.1 percent for job stayers in column (4).[15]

We also find that the real weekly earnings (excl. overtime) of new hires decline by 4.6 percent if the unemployment rate increases by one p.p., while job-stayer earnings decline by 2.9 percent. 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 Equation (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

[12] Online Appendix Tables C3 & C4 display the underlying values of all series in Figure 1.

[13] A modified Breusch-Pagan test, where we regress the squared residuals from the second stage regression, Equation (2), on the inverse of the number of jobs per period, returns non-significant coefficient estimates (p-value>0.1).

[14] A Ramsey RESET test could not reject the hypothesis that non-linear combinations of the explanatory variables have no additional explanatory power. Particularly, higher-order unemployment rate terms do not have significant additional explanatory power.

[15] We also check the effect of making 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 absolute semi-elasticity of hiring hours significantly increases in the number of minimum hires (see column (3) of Online Appendix Table B2), nearly doubling when we require at least ten 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 and job stayers to six per job and year. However, the main findings are unchanged: real wages of new hires are equally responsive, and hiring hours are more responsive to the business cycle.

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	-2.56^{***}	-2.60^{**}	-1.71^{***}	-0.10
	(0.94)	(1.13)	(0.47)	(0.18)
2. Including controls for share of full-time workers	-2.48^{***}	-2.59^{**}	-0.44^{**}	-0.11
	(0.92)	(1.12)	(0.19)	(0.14)
3. Job hires in at least 25% of years when firm is observed	-2.29^{**}	-2.61^{**}	-0.65^{**}	-0.16
	(0.94)	(1.11)	(0.31)	(0.13)
4. All jobs observed in at least 2 years	-2.36^{**}	-2.90^{**}	-0.52^{**}	-0.16
	(0.98)	(1.15)	(0.21)	(0.09)
5. Baseline sample, but weighted by number of employees per year	-2.14^{***}	-1.88	-2.74^{***}	-0.43
	(0.75)	(1.03)	(0.85)	(0.22)
6. Including other pay and hours	-2.65^{***}	-2.61^{**}	-1.64^{***}	-0.34^{***}
	(0.91)	(1.10)	(0.40)	(0.12)
7. 3-digit occupations (jobs)	-2.28^{***}	-2.65^{***}	-1.08^{***}	-0.22^{**}
	(0.79)	(0.89)	(0.40)	(0.09)

Notes.- second-step regression results, estimates $\hat{\gamma}$: responses of the period-fixed effects $\widehat{\beta}_t$ to the unemployment rate; regression specifications as in Equations (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. 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. Seventh row identifies a sample of jobs at the 3-digit level within firms, rather than the 4-digit level, and is otherwise comparable with row 1.
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.

hours the semi-elasticity estimate then increases to -0.4: over three-quarters 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 percent. We find a slightly smaller response of hiring wages to the recession, while the estimated response of hours decreases to less than half of the baseline estimate, though remains statistically significant. This suggests that not keeping the sample of jobs fixed induces problems of endogenous sample selection, since jobs with relatively rigid hiring wages and hours stopped hiring during the recession.

As previously explained, using 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 resemble the series from the baseline sample (Online Appendix Figures D4 and Figures D5). 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 in the baseline, and the fourth row in Table 3 shows comparable estimates to the baseline. In this more UK-representative sample, the real wages of all new hires and job stayers decrease by 2.4 and 2.9 percent, respectively, for each p.p. increase in the unemployment rate. However, the estimated response of hiring hours is less pronounced, though still significant at the five percent level. Thus, jobs with less flexible hiring wages and hours were more likely to stop hiring altogether during the Great Recession.

The fifth row of Table 3 shows values for $\widehat{\gamma}$ when we estimate the first step using WLS, with weights proportional to the number of employees in each job. These estimates 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 potential selection issues in other studies which use worker-level data (e.g. Carneiro et al., 2012; Haefke et al., 2013; Gertler et al., 2016). The semi-elasticity estimate for hiring wages of -2.1 is the largest of all the specifications described here. Therefore, weighting jobs by the number of hires reduces the estimated response of real wages to business cycle conditions: relatively more hires are made in jobs with robust (or rigid) hiring wages. Jobs which hired 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 60 percent larger than without endogenous weighting. Therefore, weighting by hiring volume overestimates the responsiveness of hiring hours within jobs. Perhaps unsurprisingly, the firms who hire relatively more during recessions appear to be able to move their workforce toward shorter hours and part-time working.[16] 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 the bias induced by the endogenous hiring volume of jobs, which is more in line with the hypotheses by Gertler and Trigari (2009) and Martins et al.

[16]These findings are not driven by the recursive updating of firm-identifiers and the associated lower number of new hires prior to 2003. When we repeat the estimations using only data from 2003-2016, we find that weighting by the number of new hires and job stayers gives the same qualitative responses compared to non-weighted estimates; hiring wages are less cyclical (-2.48) and hiring hours are more cyclical (-1.74). We thank an anonymous reviewer for making us aware of this potential concern.

(2012): the German data only provide information on annual earnings and the number of days worked, and so Stüber's measures of real wages are better described as average daily earnings. As our estimates show, the increase in the estimated response of hiring hours is large and could exceed the decline in the estimate for wages. When combined this could create the impression that earnings are more cyclical when weighted by hiring volume, while wage rates are in fact less cyclical.

The sixth row of Table 3 shows estimates including 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 overtime hours are more responsive than standard working hours.

The final row of Table 3 shows estimates where we identify jobs using an employee's 3-digit occupation code within a firm, instead of using their 4-digit code. Otherwise, these estimates are obtained in the same way as the main results in row 1. The number of new hires observations in this sample is approximately 20 percent larger; the number of entry-level jobs increases to 444 from 391 and the number of firms increases to 382. The main results for wages are qualitatively unchanged as we increase the representatives of the sample in this way, reducing the median firm size by over 3,000 employees (42%), with the hiring and job-stayer hourly real wage responses still not significantly different. The estimated hiring hours semi-elasticity is increased to -1.1, though this is still significantly greater in magnitude than for job stayers. However, we prefer the baseline results because we can better address the changing composition of hiring over the business cycle when jobs are more precisely defined.

3.2 Unobserved heterogeneity of new hires and jobs over the cycle

The unobserved heterogeneity of new hires and jobs over the cycle is a primary concern here. For example, if the productivity of new hires is procyclical, then low wages when unemployment is high just reflect low worker quality, such that real hiring wages per *efficiency unit* of labor might be acyclical or even countercyclical. Similarly, if a high unemployment rate forces firms with relatively rigid real wages out of the market, then the composition of the jobs sample becomes endogenous.

To check whether the main results are affected by the composition of jobs over the cycle, we create a balanced panel by including only those jobs which fulfill the selection criteria throughout 2003-16. The sample size of new hires is much smaller, with only 10,542 observations, which is around 20% of the hires in the baseline sample. In this case, we find statistically significant semi-elasticities for hourly hiring wages of -2.5 and for hours worked of -1.3. These estimates indicate that the composition of jobs has somewhat changed over the cycle. Specifically, the smaller responses suggest the jobs that entered (left) the sample exhibit a slightly larger (smaller) degree of flexibility in real hiring wages and hours. But, these estimates are not substantially different from the baseline results, thus suggesting that the changing jobs composition is not a large problem for our approach.

We also test for possible bias resulting from the selection of jobs from the ASHE panel dataset into the baseline sample of jobs, using a method suggested by Wooldridge (2010). We first create an indicator variable, for each job and year, which equals one if the job fulfills the entry-level job criteria in that year and zero otherwise. We then separately regress leads and lags of this dummy on job-fixed effects, all the covariates included in Equation (1) and the unemployment rate. The coefficient estimate on the unemployment rate is then informative on the presence of selection bias: if the estimates are not statistically significant, then the selection of jobs is uncorrelated with the business cycle, conditional on observable characteristics, and thus there is less concern about systematic differences between the jobs which are selected into the sample and those which are not. If we restrict the "population" of jobs to those in the Wholesale & Retail, Hotels & Accommodation, Non-financial Professional Services, and Recreation & Sports industries, then the estimated coefficient of the unemployment rate is not statistically significant (p-value> 0.10). The most notable omission is the financial services industry, which almost completely stopped hiring with the onset of the 2008 financial crisis. The other industries, for which the baseline sample of jobs is representative, account for around half of private sector employment and for two-thirds of private sector hires in the UK.

To address the other concern of cyclicality in the quality of new hires, we would ideally use a proxy for unobserved quality, such as employees' highest qualifications. Unfortunately, the ASHE does not contain information on education or qualification levels. In the previous section, we discussed how some relevant observable characteristics of new hires do not show any notable cyclical pattern. Here we provide further evidence that the main results are not driven by the composition of the workforce, using two different approaches.

First, we propose a simple test: we construct a "quality indicator" by computing the job-level median, within the baseline sample jobs, of new hires' last observed real wages at their previous employer. We then regress this variable on the unemployment rate, following exactly the same two-step approach as before. Intuitively, if real wages and productivity are correlated at the worker level, finding a significant coefficient for the unemployment rate in this regression would suggest that the quality of new hires systematically varies over the business cycle. However, if real wages are not a good proxy for labor productivity, then not controlling for the unobserved heterogeneity of new hires should be of less concern for the estimates presented in the previous section. The estimated semi-elasticity of this worker quality indicator with respect to unemployment is -0.43, with a standard error of 1.15, which is not significant at any of the usual levels. A potential problem with this test is that the Great Recession affected employment in the financial services industry more negatively than elsewhere; high wage-premiums in this industry might increase the quality indicator during the recession regardless of unobserved worker productivity. Therefore, as a further check, we repeat the two-step regression approach described in Section 2 using an adjusted measure of worker quality. Specifically, the dependent variable in the first step is now the median within a job of the previously observed wage of new hires minus the corresponding previous industry's mean wage. By subtracting the industry mean, we control for industry-specific wage premiums. The estimates of the semi-elasticity with respect to the unemployment rate are all negative and not statistically significant, regardless of the industry classification level used.[17]

Second, we assess the potential presence of cyclical and wage-relevant unobserved worker heterogeneity among new hires using the UK Labour Force Survey (LFS). The LFS is a quarterly household-based survey, representative of the UK population. We follow Blundell et al. (2014) and define skill-levels for three categories of a worker's highest qualification within the LFS: less than 5 GCSEs (General Certificate of Secondary Education) at grades A*-C or equivalent is classified as low skill; a university degree or equivalent is classified as high skill; and intermediate qualifications are classified as medium skill. Individuals in the LFS are asked how many months they have been working for their current employer, and we define a new hire if a worker was hired between May of the previous year and April of the current year, corresponding to the observation period in the ASHE. To measure the cyclicality of hiring by skill-level, we construct an indicator that equals one if the employee is a new hire and zero otherwise. We estimate a linear probability model, again using a two-step approach to deal with clustering at the annual level. In the first step, we separately regress the indicator for each skill level on time-period dummies. The coefficient estimates of these dummies are then regressed on a linear trend and the unemployment rate. The coefficient of the unemployment rate shows the marginal effect of a one percentage point increase of the unemployment on the probability of hiring. The results are presented in Table 4, where we include only new hires in industries that represent the ASHE baseline sample of jobs. The first row shows that workers of all skill levels are less likely to be hired when the unemployment rate increases, but only the coefficients for low- and medium-skilled

[17] The estimates (standard errors) are: -0.54 (0.49) at the 1-digit level; -0.84 (0.53) at the 2-digit level; -0.69 (0.64) at the 3-digit level; and -0.72 (0.66) at the 4-digit level.

workers are statistically significant. When the unemployment rate increases by one percentage point, the probability that a low- or medium-killed worker is hired decreases by 23 or 12 percentage points, respectively. In the second row of Table 4, we exclude hires in financial services from the first-step regression. The results are almost unchanged in this case.

TABLE 4: Estimated marginal effects of the unemployment rate on the probability of hiring, linear probability model, 1998-2016.

	Hiring probability (hire=1, 0 otherwise)		
	Low-skill (1)	Medium-skill (2)	High-skill (3)
1. Included sectors: SIC2007: 47, 55, 56, 64, 78, 81	-0.23^{***} (0.04)	-0.12^{**} (0.06)	-0.11 (0.06)
2. Included sectors: SIC2007: 47, 55, 56, 78, 81	-0.23^{***} (0.04)	-0.14^{**} (0.06)	-0.13 (0.06)

Notes.- source: UK Quarterly Labour Force Survey. SIC2007 refers to Standard Industrial Classification (2007) categories: (47) Retail trade, excl. motor vehicles and motorcycles; (55) Hotels & (56) Restaurants; (64) Financial service activities, excl. insurance and pension funding; (78) Employment activities; (81) Services to buildings and landscape activities.
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.

To test whether the hiring probabilities respond differently to the business cycle, we regress the differences between skill levels of the first-step time-period dummy estimates on the unemployment rate. Only the cyclicality of medium-skilled hiring probability differs significantly from the low-skilled hiring probability at the 5% level. The difference is positive, indicating that low-skilled workers are relatively less likely to be hired compared with medium-skill workers. Though this analysis does not control for the composition within jobs, these results suggest that the unobserved quality of new hires in the UK was countercyclical over the period 1998-2016 and within the particular set of industries studied. This implies that our main estimates for the size of the responses of real hiring wages and hours to the business cycle could be lower bounds if interpreted in terms of efficiency units of labor.

To summarize, the findings in this subsection suggest that the quality of the new hires among entry-level jobs is not substantially varying over the business cycle. This reassures us that unobserved changes in the composition of jobs and the employees within them are not a first-order concern for the main results, although it is not possible to be definitive on this point.

4 Further robustness and discussion

The main results described above show that UK firms were able to significantly decrease the real labor cost per employee in response to the Great Recession. To further address robustness, in this section we apply some alternative estimation procedures. All 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 5 displays results from varying the specification of the second-step regression (2), while leaving the first step 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 the dataset.

TABLE 5: Estimated semi-elasticity of real wages and hours with respect to the unemployment rate, 1998-2016: alternative specifications of the second-step regression

	Wages		Hours	
	New hires (1)	Job stayers (2)	New hires (3)	Job stayers (4)
1. Baseline (OLS)	-2.56^{***}	-2.60^{**}	-1.71^{***}	-0.10
	(0.94)	(1.13)	(0.47)	(0.18)
2. Baseline with quadratic trend	-2.50^{***}	-1.96^{***}	-1.69^{***}	-0.13
	(0.43)	(0.42)	(0.49)	(0.13)
3. Baseline sample, but weighted by number of jobs per year	-2.57^{***}	-1.90^{**}	-1.72^{***}	-0.44^{**}
	(0.85)	(0.88)	(0.51)	(0.19)
4. First differences (OLS)	-1.61^{***}	-1.84^{***}	-0.11	-0.09
	(0.61)	(0.40)	(0.61)	(0.16)

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 applies WLS in the second step, with weights in proportion to the number of jobs observed per year. Fourth row estimates Equation (2) in first differences. 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.

For comparability with Martins et al. (2012), we re-estimate Equation (2) using WLS, with weights proportional to the number of jobs per period in the first step. The resulting estimates in the third row of Table 5 are qualitatively unchanged from the baseline. However, the real wages of job stayers are slightly less cyclical. In the final row of Table 5, we also re-estimate Equation (2) in first differences, to address potentially spurious estimates if wages, hours, or the unemployment rate are integrated. As in the 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 unemployment rate increases by one p.p., then hiring wages decrease by 1.6 percent, and for job-stayers' wages decrease by 1.8 percent. This is comparable to the finding by Devereux and Hart (2006) of 1.7 percent for all job stayers in the UK during 1975-2001.

Overall, the results that both the real wages of hires in entry-level jobs and of job stayers declined in response to increases in the unemployment rate 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 first-difference version of Equation (2), which indicates that the decrease in hiring hours is better understood as a medium-run and persistent development since 2008. An issue we cannot robustly address here, given the sample period we study includes only one official recession, is whether the wage response to the business cycle is symmetric. Font et al. (2015) found some evidence from Spain using a panel of workers that real wages were less responsive in recessions than in expansions, though these asymmetric effects were less pronounced among new hires than among job stayers. In Online Appendix B, we discuss the results of several more robustness checks, which also do not affect our confidence in the main results.

In Online Appendix F, we further analyze the impact that the UK's National Minimum Wage could have on our estimates. Using the partial equilibrium set up of DiNardo et al. (1996), we compute hiring wages under the counterfactual assumption that the minimum wage is not constraining firms in their ability to decrease hiring wages during the Great Recession. We estimate that the absolute value of the semi-elasticity of real hourly hiring wages to the unemployment rate would have been around 0.2 percentage points larger in the absence of the National Minimum Wage.

4.2 Using labor productivity as the business cycle indicator

As an alternative indicator for the business cycle we consider labor productivity, measured by log real gross value added per hour.[18] Measures of labor productivity are particularly relevant to 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 labor market or other theories of business cycles: if real wages are perfectly rigid, then they should not respond to labor productivity. Table 6 shows the estimated elasticity of real wages and hours with respect to labor productivity, when we replace the unemployment rate with labor productivity in regression Equation (2). In the first row we use aggregate labor productivity as the business cycle indicator. The estimates are smaller than one, but positive

[18] Source: ONS Labour Market Statistics, April 2017, available at https://www.ons.gov.uk/.../apr2017; accessed 24/04/2017. Online Appendix Figure D4 shows the time series of each business cycle indicator used and Online Appendix Table C5 shows the underlying values.

and statistically different from zero. Hiring hours also respond significantly to labor productivity, though less than real wages do.

TABLE 6: Estimated elasticity of real wages and hours with respect to labor productivity, 1998-2016

	Wages		Hours	
	New hires (1)	Job stayers (2)	New hires (3)	Job stayers (4)
1. Whole economy	0.78^{***}	0.94^{***}	0.28^{***}	-0.01
	(0.10)	(0.07)	(0.10)	(0.04)
2. Services sector	0.84^{***}	1.03^{***}	0.30^{***}	-0.01
	(0.12)	(0.09)	(0.11)	(0.04)

Notes.- second-step regressions of estimated period effects on alternative indicators of the business cycle. First-step estimated according to (1) and (3). First row uses the log of real whole economy gross value added (GVA) per hour: ONS series LZVB. Second row uses the log of real gross value added (GVA) per hour in Services (SIC 2007: 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 percent of jobs in the baseline sample belong to the services industry, and the response of labor productivity to the Great Recession differed markedly across sectors, we also use services sector labor productivity as the cyclical indicator. The second row of Table 6 then shows a higher estimated elasticity of real wages and hours worked. The real wages of new hires and job stayers significantly decrease by 0.8 and one percent, respectively, when aggregate labor productivity decreases by one percent. The difference between these values is insignificant. The estimates for job stayers and new hires do not significantly differ from one. Given we would expect no relationship between aggregate productivity and real wages if the latter were perfectly rigid, there is evidence here to reject this hypothesis for both new hires and job stayers in the UK.

The estimated hiring wage elasticity here with respect to aggregate productivity 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 the average real daily earnings of incumbent German workers increase by 0.5 percent if aggregate real output per worker increases by one percent, and he estimates a significantly smaller coefficient for new hires.

5 How do wages and hours evolve after hiring?

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 labor 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 profit is potentially less cyclical than measured for the hiring wage. In this way, 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 in labor costs, Figure 2 plots the real hourly wages and hours worked, averaged over employees instead of jobs, for each cohort of entry-level job 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. Figure 2A suggests that wages exhibit cohort effects: the real wages of hiring cohorts from 1998 to 2005 generally 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 are locked into low-wage growth trajectories. For example, the mean wages of the 2013 cohort of hires in 2015 were below the mean wages of the 2014 cohort in 2015. Figure 2B 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 labor costs.

Comparing sample averages over time is likely to be subject to a composition bias, since relatively low-wage employees are less likely to be employed and so 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 only look at 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

FIGURE 2: Paths of real wages and hours for cohorts of new hires

A. Mean hourly wages

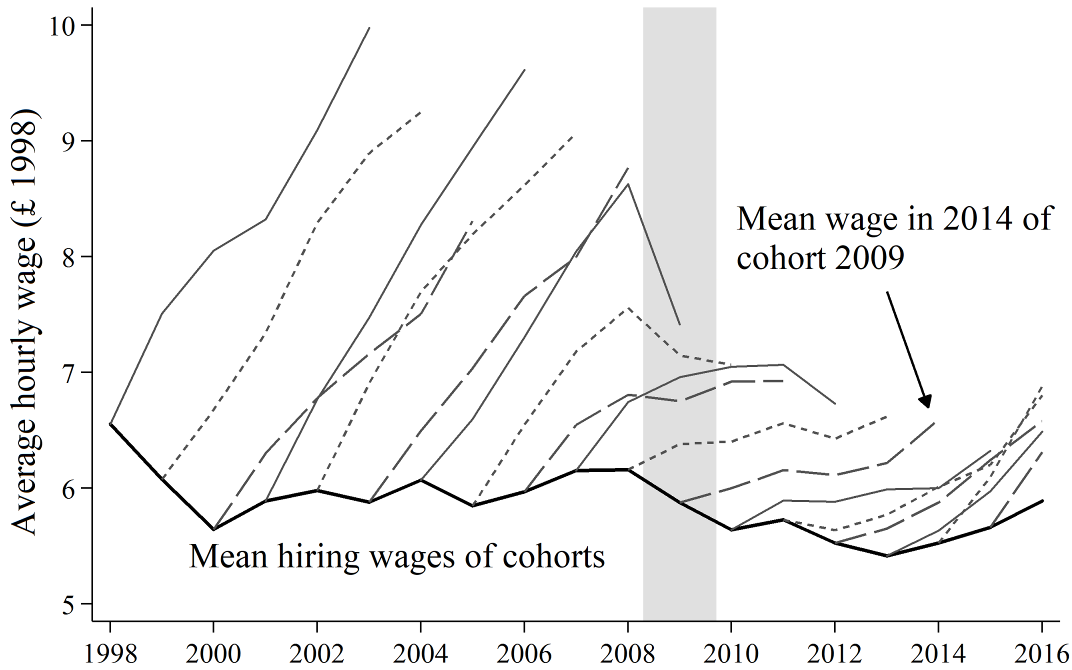

B. Mean hours worked

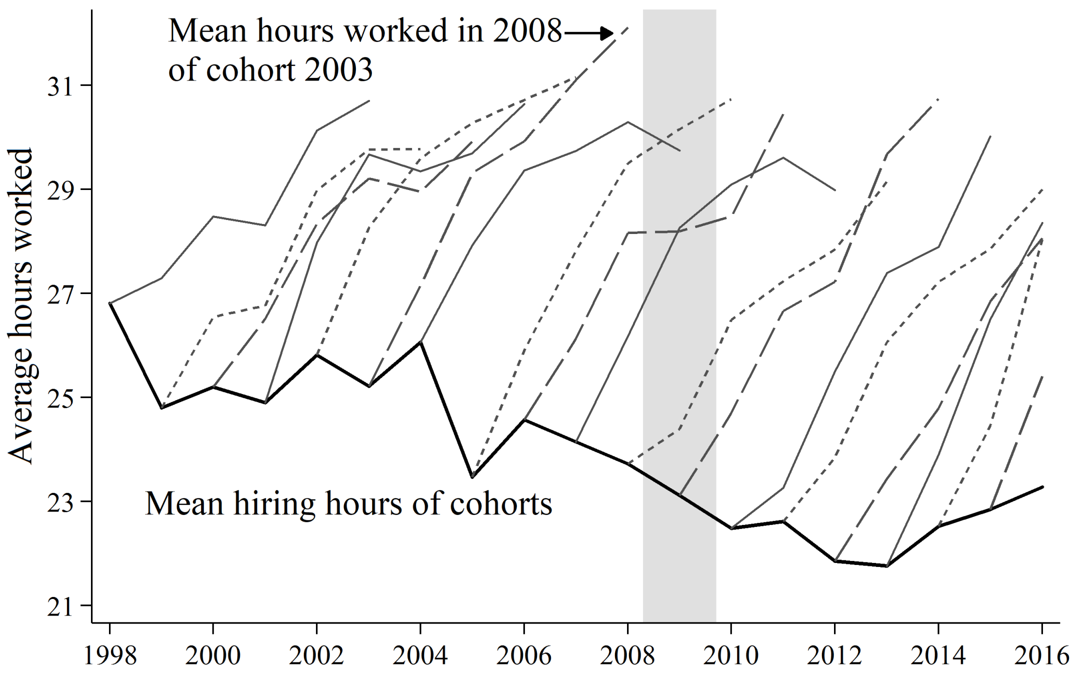

Notes.- the solid lines give the average real hiring wage and weekly hours worked for each cohort of new hires in the sample of entry-level jobs (i.e. column (1), Table 1). Each line branching off from the solid line shows the paths of wages or hours of these hiring cohorts over time, as their tenure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

firms, or they exit into non-employment. Using least squares, we estimate:

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

where the dependent variable is the log real wage of some match m between a worker and job with tenure τ. θ_m is a match-fixed effect and $\psi_{c\tau}$ are cohort-tenure-dummies for any matches beginning in year $c = 2001,...,2013$ 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, since employees leave jobs over time. 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 Equation (4) by excluding the dummies ψ_{c0}, so the estimated values of $\psi_{c\tau}$ for $\tau > 0$ are interpreted as average 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 Equation (4) with log basic weekly hours worked as the dependent variable.[19]

We plot the estimated cohort-tenure-dummies in Figure 3 for selected hiring years, and the underlying estimates and standard errors are displayed in Online Appendix Table C6 for all years. For example, the three left-most bars display the wage growth of the 2005 cohort of hires, the last cohort unaffected by the Great Recession within three completed years of tenure. The first bar and the corresponding entry in Online Appendix Table C6 (row "2005", column "Wages, 1 year") show that workers who were hired in 2005 experienced average annual real hourly wage growth of five percent over the following year, conditional on still working in the same job in 2006. Those workers of the 2005 cohort who were still working in the same job in 2008 saw their wages increase by nine percent compared to their hiring wage, as the third bar shows. Figure 3A 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 percent. Similarly, there was no real wage growth over the subsequent three years for 2009 hires, compared with over eight percent for 2005 hires. These results are statistically significant.

The cyclical differences in hours growth in these jobs are less pronounced in Figure 3B. 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

[19] The sample of new hires used for this analysis is a subset of the 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 percent smaller. The estimated real wage semi-elasticity with respect to the unemployment rate for new hires is 3.2 and is slightly larger in absolute terms than in the baseline sample, while hours worked are just as responsive as measured before.

FIGURE 3: Estimated composition-adjusted and cohort-specific log changes $\widehat{\psi}$ in real wages and hours relative to hiring levels: workers who stay in entry-level jobs

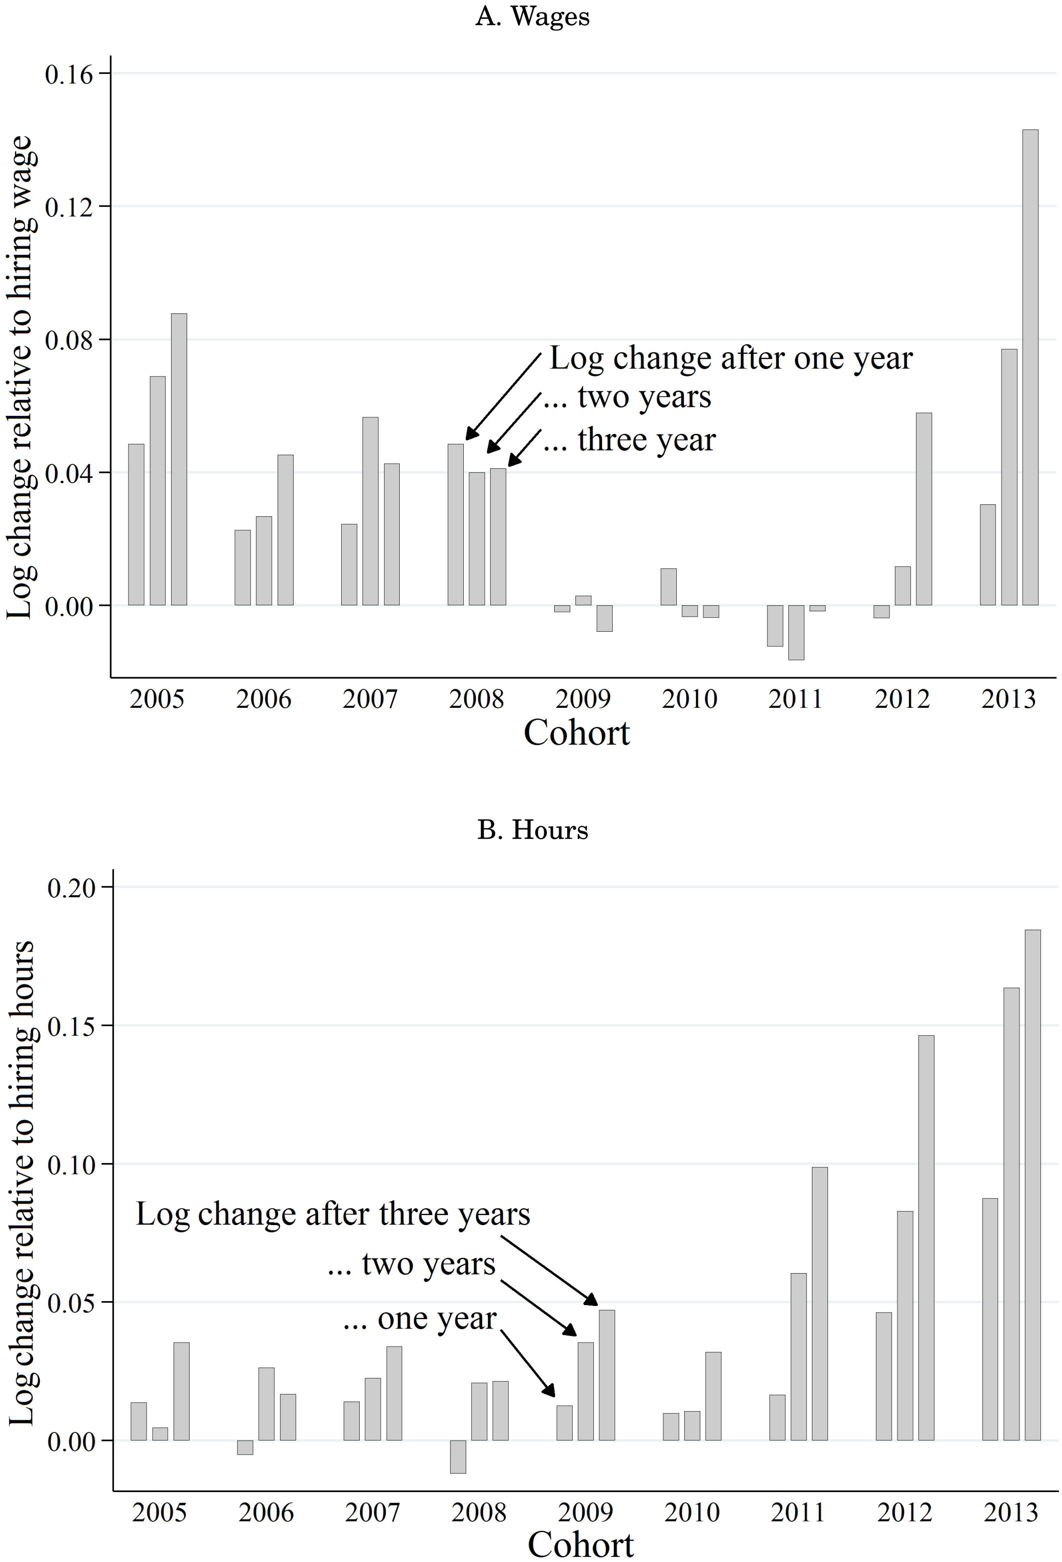

Notes.- cohort average change in log wages and log hours with tenure, relative to respective hiring values, in entry-level jobs. Composition adjusted by controlling for match-fixed effects. See Online Appendix Table C6 for standard errors and results for all other hiring years in 2001-13.

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. We find that nearly two-thirds of the changes in hours worked within jobs are due to workers switching between part- and full-time hours status, despite the fact that only 10-15 percent of workers actually change their status. Specifically, Online Appendix Tables C7-C9 display descriptive statistics on the transition rates between full-time and part-time employment, and vice versa, and the associated average changes in hours worked, as well as the share of the overall change in average hours worked explained by switching status. Within the first year on-the-job, the likelihood of a new hire switching to full-time hours within a match is around 60 percent larger than the likelihood of switching from full-time to part-time. The likelihood of switching generally falls with tenure, but remains one-third higher for switches from part-time to full-time than vice versa. The post-recession period of 2008-13 exhibits higher propensities for switching hours status within job than the pre-recession period of 2002-7.

Thus, the pattern within these particular jobs suggests a refinement to the findings of Borowczyk-Martins and Lalé (2019), who use worker-level flows to show that within-firm transitions from full-time to part-time mostly accounted for the rise and persistence of part-time employment during the UK's Great Recession. We have found that at the job-level the hours worked by new hires are especially cyclical compared with job stayers. These new hires then experience an elevated rate of hours changes during recessions. Additionally, the findings of Borowczyk-Martins and Lalé apply to higher frequency data (they use the longitudinal version of the UK Quarterly Labour Force Survey), whereas we can confirm similar patterns at an annual frequency for new hires. Similarly, Kurmann and McEntarfer (2018) document that US employees who stay for at least two years in the same firm, as opposed to in the same job, experience cyclical variation in hours worked. Future research should try to address whether 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, and what persistent effects the labor market state at the time of hiring has.

It is worth noting that the procyclical hours growth for new hires who remain in entry-level jobs can somewhat compensate for the procyclical response of initial hiring hours. However, this on-the-job hours growth only sufficiently compensates the hiring cohorts in 2012 and 2013, after being employed for three years, for the estimated decrease in initial hours of around 15 log points (see Section 3). Despite this, the sluggish real wage growth estimated within jobs means the weekly *earnings* of new hires, and thus labor costs, remain substantially procyclical.

The findings in this section suggest that firms were not only able to significantly reduce the real wages and hours worked of new hires' hours 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 of Equation (4) is subject to the same measurement criticism that 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 here due to a small number of degrees of freedom at the job level.

6 Concluding remarks

We provide new estimates on the flexibility of UK wages. Most importantly this is measured at the job level, which is appropriate for understanding how firms adjust their per employee labor costs in response to business cycle conditions in frictional labor markets. We find that job-stayer real wages respond by as much as 2.6 percent for a one percentage point rise in the unemployment rate. The elasticity of job-stayer real wages with respect to aggregate labor productivity equals approximately one. Hiring wages are as responsive to the business cycle as the wages of job stayers. This conforms with results from other countries, suggesting that rigid real hiring wages are not the appropriate way to model and understand 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 remains as to why firms were able to adjust wages so freely, and why workers were so willing to accept these changes. In Online Appendix G, we speculate about some possible UK-specific explanations for this magnitude of measured wage flexibility.

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 unemployment rate. Conversely, the hours of new hires among the same firms responded significantly, decreasing by 1.7 percent for a 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 behavior, which should in the first instance be tested outside the specific UK context, and subsequently reflected on when modeling how firms adjust their workforce to shocks.

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 labor costs from the firm's perspective for hires made during this time. In this respect, it seems that firms experienced even more flexibility 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 labor costs per new employee.

References

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

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

Bils, M. J. 1985. "Real Wages over the Business Cycle: Evidence from Panel Data." *Journal of Political Economy*, 93(4): 666–689.

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.

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

Borowczyk-Martins, D., and E. Lalé. 2019. "Employment Adjustment and Part-Time Work: Lessons from the United States and the United Kingdom." *American Economic Journal: Macroeconomics*, 11(1): 389–435.

Brown, W., P. Marginson, and J. Walsh. 2003. "The management of pay as the influence of collective bargaining diminishes." Blackwell Publishing, 189–212.

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.

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.

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

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.

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.

Donald, S. G., and K. Lang. 2007. "Inference with Difference-in-Differences and Other Panel Data." *Review of Economics and Statistics*, 89(2): 221–233.

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.

Font, P., M. Izquierdo, and S. Puente. 2015. "Real wage responsiveness to unemployment in Spain: asymmetries along the business cycle." *IZA Journal of European Labor Studies*, 4(1): 1–13.

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.

Gertler, M., and A. Trigari. 2009. "Unemployment Fluctuations with Staggered Nash Wage Bargaining." *Journal of Political Economy*, 117(1): 38–86.

Goos, M., and A. Manning. 2007. "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain." *Review of Economics and Statistics*, 89(1): 118–133.

Gregg, P., S. Machin, and M. Fernández-Salgado. 2014. "Real Wages and Unemployment in the Big Squeeze." *The Economic Journal*, 124(576): 408–432.

Haefke, C., M. Sonntag, and T. van Rens. 2013. "Wage rigidity and job creation." *Journal of Monetary Economics*, 60(8): 887–899.

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).

Kudlyak, M. 2014. "The cyclicality of the user cost of labor." *Journal of Monetary Economics*, 68(C): 53–67.

Kurmann, A., and E. McEntarfer. 2018. "Downward Wage Rigidity in the United States: New Evidence from Administrative Data." mimeo.

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.

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

Office for National Statistics. 2017. “Annual Survey of Hours and Earnings, 1997-2016: Secure Access. [data collection].” 9th Edition SN: 6689, UK Data Service.

Office for National Statistics. 2018. “Quarterly Labour Force Survey, 1992-2018: Secure Access. [data collection].” 14th Edition SN: 6727, UK Data Service.

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

Salvatori, A. 2018. “The anatomy of job polarisation in the UK.” *Journal for Labour Market Research*, 52(1): 1–15.

Shin, D. 1994. “Cyclicality of Real Wages among Young Men.” *Economics Letters*, 46(2): 137–142.

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.

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

Stüber, H. 2017. “The Real Wage Cyclicality of Newly Hired and Incumbent Workers in Germany.” *The Economic Journal*, 127(600): 522–546.

Wooldridge, J. M. 2010. *Econometric Analysis of Cross Section and Panel Data*. MIT Press.

Cyclical Labor Costs Within Jobs[†]
Online Appendix

Daniel Schaefer[‡] Carl Singleton[§]

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 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 the 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 random sample of all employees, irrespective of 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 percent response, the ASHE is a true one percent random sample of employees: all those 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 small, we do not believe they could affect the 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 the 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

[†]This work 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.

[‡]School of Economics, University of Edinburgh; daniel.schaefer@ed.ac.uk

[§]Department of Economics, University of Reading; c.a.singleton@reading.ac.uk

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." (ONS). From 2007, the sample size of the ASHE was reduced by 20 percent, 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. After 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 2003-16 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 the 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 rates of pay, excluding overtime, are less than 80 percent 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 behavior which has been hypothesized by the ONS. Since the data are not top-coded, we drop the highest one percent of all weekly or hourly earners.

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 started working for an enterprise from 2002 onward. 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 the CH-firms sample of new hires if we are in any doubt. Given the main 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 the employment start dates 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 a year and an 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. It causes a substantial additional degree of polarization of work from 2002 onward. 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 dual 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 the baseline/main result for convenience. The robustness discussed here is with regards to the specification of the first step of the regression models: (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 of wages and hours are employee sample means, rather than median values. Qualitatively the results are unchanged: wages for hires and job stayers exhibit a sizable and significant cyclical response, as do hiring hours. 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 the specific sample selection criteria for jobs. However, for robustness, rows 3 and 4 show estimates using the 25th and 75th percentiles of within job wages or hours as the first-step dependent variables, respectively. For wages, the responses are qualitatively unchanged, though the estimated response to the unemployment rate at the 75th percentile is slightly higher, reflecting that these hires were less constrained cyclically by the National Minimum Wage. The estimated hours response for job stayers remains small in magnitude at both percentiles within jobs, though it is statistically significant from zero at the 75th percentile. The hiring hours response was insignificant for hires with relatively low hours within jobs, i.e. part-time. However, at the 75th percentile within jobs, the hiring hours response to the business cycle is greater than at the median. The fifth 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 sixth 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 unemployment rate than in the private sector, potentially reflecting the squeeze on labor costs imposed by fiscal austerity. The seventh 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.

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	-2.56^{***}	-2.60^{**}	-1.71^{***}	-0.10
	(0.94)	(1.13)	(0.47)	(0.18)
2. Job means	-2.65^{***}	-2.67^{**}	-1.21^{***}	-0.28
	(0.96)	(1.17)	(0.39)	(0.18)
3. 25th percentile	-1.98^{**}	-2.69^{**}	-0.21	-0.35
	(0.82)	(1.23)	(0.53)	(0.43)
4. 75th percentile	-3.27^{***}	-2.76^{**}	-2.24^{**}	-0.42^{***}
	(1.11)	(1.16)	(0.99)	(0.14)
5. Baseline, but without controls	-2.56^{***}	-2.61^{**}	-1.87^{***}	0.02
	(0.86)	(1.10)	(0.41)	(0.21)
6. Baseline, but including public sector	-2.32^{**}	-2.62^{**}	-2.07^{***}	-0.41^{***}
	(1.01)	(1.31)	(0.33)	(0.10)
7. RPI instead of CPI	-1.96^{***}	-1.97^{**}		
	(0.71)	(0.90)		

Notes.- second-step regression results of estimated period effects on the unemployment rate, $\widehat{\gamma}$. First row is identical to Table 3, included here for comparison. Second row uses mean wages or hours in jobs as the dependent variable in the first step. Third row uses the 25th percentiles of wages and hours within jobs as the dependent variable in the first step. Fourth row uses the 75th percentiles of wages and hours within jobs as the dependent variable in the first step. Fifth row excludes all time-varying controls from the first step. Sixth row includes public sector firms in the analysis. Seventh row uses the Retail Price Index, instead of the Consumer Price Index, to deflate wages.

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 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	-2.42^{***}	-2.66^{**}	-1.18^{***}	-0.26^{**}
	(0.83)	(0.94)	(0.35)	(0.11)
3 (baseline)	-2.56^{***}	-2.60^{**}	-1.71^{***}	-0.10
	(0.94)	(1.13)	(0.47)	(0.18)
4	-2.63^{***}	-2.46^{***}	-2.08^{***}	-0.11
	(0.83)	(0.91)	(0.54)	(0.15)
5	-2.84^{***}	-2.22^{**}	-1.97^{***}	0.01
	(0.80)	(0.89)	(0.55)	(0.17)
6	-2.88^{***}	-2.15^{***}	-2.21^{***}	-0.28
	(0.77)	(0.80)	(0.67)	(0.17)
7	-2.85^{***}	-2.25^{**}	-2.23^{***}	-0.38
	(0.78)	(0.88)	(0.58)	(0.21)
8	-2.72^{***}	-2.27^{**}	-2.04^{***}	-0.56^{***}
	(0.73)	(0.94)	(0.77)	(0.21)
9	-2.70^{***}	-2.41^{***}	-2.66^{***}	-0.55
	(0.70)	(0.93)	(0.90)	(0.23)
10	-2.81^{***}	-2.01^{***}	-2.61^{***}	-0.49
	(0.63)	(0.75)	(0.95)	(0.25)

Notes.- second-step regression results of estimated period effects on the unemployment rate, $\widehat{\gamma}$. 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 row 1, 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 $\widehat{\beta}_t$ series between new hires and job stayers with respect to the unemployment rate/labor productivity, 1998-2016

	Wages	Hours
Baseline sample	-0.35	-1.57^{***}
	(0.31)	(0.43)
Table 3		
2. Including controls for share of full-time workers	-0.27	-0.30
	(0.32)	(0.21)
3. Job hires in at least 25% of years when the firm is observed	-0.02	-0.48
	(0.33)	(0.33)
4. All jobs observed in at least 2 years	0.09	-0.36
	(0.28)	(0.25)
5. Baseline sample, but weighted by number of employees per year	-0.59	-2.23^{**}
	(0.39)	(0.98)
6. Including other pay and hours	-0.44	-1.28^{***}
	(0.30)	(0.41)
7. 3-digit occupations (jobs)	0.04	-0.87^{***}
	(0.35)	(0.31)
Table 5		
2. Baseline with quadratic trend	-0.37	-1.55^{***}
	(0.21)	(0.35)
3. Baseline sample, but weighted by number of jobs per years	-1.00^{***}	-1.20^{**}
	(0.29)	(0.50)
4. First differences (OLS)	0.25	-0.04
	(0.36)	(0.61)
Table 6		
1. Whole economy	-0.12^{**}	0.29^{***}
	(0.06)	(0.10)
2. Services sector	-0.14^{**}	0.31^{***}
	(0.06)	(0.11)

Notes.- "second-step" regression results of estimated period effects on the unemployment rate/labor productivity, $\widehat{\gamma}$, as per (2), where the Dependent variable is instead the difference between hiring and job stayer composition-adjusted period means, $[\{\widehat{\beta}_t\}_{\text{New hires}} - \{\widehat{\beta}_t\}_{\text{Job stayers}}]$, obtained from associated first-step regressions described by the corresponding rows in Tables 3, 5 and 6.

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
Retail Trade; Repair of Personal and Household Goods (52)	26,792	0.49
Hotels and Restaurants (55)	9,842	0.18
Financial Intermediation, exc. Insurance and Pension Funding (65)	3,821	0.06
Other Business Activities (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 ($\widehat{\beta}_t$ from first-step regressions)

	CH-firms		ASHE	
	New hires	Job stayers	New hires	Job stayers
Year	(1)	(2)	(3)	(4)
1999	-0.093^{***}	-0.091^{***}	-0.086^{***}	-0.103^{***}
2000	-0.091^{***}	-0.067^{***}	-0.084^{***}	-0.072^{***}
2001	-0.049	-0.048^{***}	-0.046^{***}	-0.039^{***}
2002	-0.012	-0.011	-0.019	-0.013^{***}
2003	0.000	0.000	0.000	0.000
2004	0.007	0.003	0.008	0.010^{**}
2005	0.025	0.026	0.028^{**}	0.033^{***}
2006	0.022	0.045	0.036^{***}	0.042^{***}
2007	0.036	0.045	0.041^{***}	0.047^{***}
2008	0.009	0.045	0.025^{**}	0.047^{***}
2009	0.031	0.048	0.026^{**}	0.034^{***}
2010	-0.001	0.017	-0.002	-0.001
2011	-0.043	-0.017	-0.027^{**}	-0.022^{***}
2012	-0.067^{**}	-0.047	-0.053^{***}	-0.050^{***}
2013	-0.073^{**}	-0.067	-0.065^{***}	-0.064^{***}
2014	-0.074^{**}	-0.072	-0.077^{***}	-0.067^{***}
2015	-0.053	-0.049	-0.057^{***}	-0.046^{***}
2016	-0.004	-0.017	-0.004	-0.020^{***}

Notes.- time series of period-fixed effects for different subsamples of the ASHE. Results based on the OLS estimations of Equations (1) and (3). Normalized to zero in 2003, linear trends removed. (1) Entry-level new hires, (2) job stayers in CH-firms (3) ASHE new hires, (4) ASHE job stayers. *** Statistically significant from zero at the 1% level; ** at the 5% level, two-sided tests, standard errors robust to clustering at the firm-level.

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

	CH-firms		ASHE	
	New hires	Job stayers	New hires	Job stayers
Year	(1)	(2)	(3)	(4)
1999	−0.003	0.007	−0.002	−0.005***
2000	0.040	0.011	0.001	−0.002
2001	0.018	0.005	0.007	−0.001
2002	0.021	−0.001	−0.002	−0.001
2003	0.000	0.000	0.000	0.000
2004	0.003	−0.002	0.023	0.001
2005	−0.010	−0.012	−0.017	−0.010**
2006	−0.027	−0.010	−0.013	−0.009**
2007	0.016	0.003	−0.012	−0.010**
2008	−0.052	0.009	−0.008	−0.008
2009	−0.053	0.002	−0.024	−0.012**
2010	−0.059	0.004	−0.037**	−0.009
2011	−0.131***	−0.004	−0.049***	−0.009
2012	−0.139***	−0.006	−0.052***	−0.009
2013	−0.154***	−0.002	−0.047***	−0.004
2014	−0.129***	0.008	−0.045***	−0.000
2015	−0.132***	0.006	−0.051***	0.003
2016	−0.119**	0.006	−0.062***	0.001

Notes.- see Table C3
*** Statistically significant from zero at the 1% level; ** at the 5% level, two-sided tests, standard errors robust to clustering at the firm-level.

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

				Labor productivity	
Year	CPI	RPI	SPPI	Whole economy	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; “Labor prod. Whole economy” - chain volume measure of gross value added at basic prices in the UK; “Labor 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		
Cohort	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
		(0.62)	(0.80)	(0.83)	(0.92)	(1.13)	(1.38)
2003		−0.09	2.63***	6.68***	0.80	2.62	2.09
		(0.74)	(0.91)	(1.01)	(1.21)	(1.49)	(1.69)
2004		2.07***	7.13***	8.61***	0.91	2.79**	0.95
		(0.61)	(0.74)	(0.89)	(1.10)	(1.26)	(1.48)
2005		4.86***	6.89***	8.78***	1.37	0.46	3.54**
		(0.50)	(0.67)	(0.71)	(1.07)	(1.38)	(1.53)
2006		2.26***	2.67***	4.53***	−0.52	2.63	1.67
		(0.49)	(0.65)	(0.79)	(1.17)	(1.54)	(1.95)
2007		2.44***	5.66***	4.27***	1.40	2.25	3.40
		(0.46)	(0.66)	(0.73)	(1.16)	(1.57)	(1.82)
2008		4.85***	4.00***	4.12***	−1.20	2.08	2.13
		(0.43)	(0.55)	(0.64)	(1.07)	(1.31)	(1.51)
2009		−0.20	0.28	−0.80	1.25	3.55***	4.72**
		(0.39)	(0.50)	(0.61)	(1.11)	(1.37)	(1.95)
2010		1.11**	−0.35	−0.37	0.98	1.06	3.20
		(0.47)	(0.61)	(0.78)	(1.40)	(1.94)	(2.01)
2011		−1.24**	−1.65**	−0.18	1.64	6.04***	9.88***
		(0.52)	(0.67)	(0.89)	(1.34)	(1.76)	(2.06)
2012		−0.39	1.17	5.79***	4.63***	8.28***	14.64***
		(0.47)	(0.65)	(0.86)	(1.20)	(1.69)	(2.05)
2013		3.04***	7.70***	14.30***	8.75***	16.35***	18.46***
		(0.53)	(0.65)	(1.00)	(1.22)	(1.72)	(2.31)

Notes.- This table reports the full results displayed in Figure 3, Section 5. “Cohort” refers to the year of hiring. Robust standard errors in parentheses.
*** Statistically significant from zero at the 1% level; ** at the 5% level, two-sided tests.

TABLE C7: Gross transition rates of job stayers by tenure between full-time and part-time work (percent of all job stayers), and associated mean change in hours worked

		PT to FT			FT to PT		
Cohort	Tenure:	1 year	2 years	3 years	1 year	2 years	3 years
2002	Pr(transition):	6.17	4.77	5.25	3.36	4.91	3.57
	Δ hours:	16.61	15.51	13.44	−15.72	−12.64	−12.27
2003		7.64	8.01	4.79	4.20	3.81	4.19
		18.04	12.75	11.87	−15.04	−14.74	−12.07
2004		8.38	6.34	3.33	3.97	4.43	3.59
		15.84	15.54	10.56	−17.27	−13.39	−12.13
2005		8.58	6.55	5.91	5.61	6.03	3.16
		16.89	13.52	15.26	−17.82	−16.28	−13.92
2006		6.69	7.75	5.71	5.94	3.01	6.29
		16.83	15.49	15.12	−20.35	−14.38	−16.43
2007		8.35	5.88	6.04	4.57	6.21	3.72
		15.20	16.45	12.16	−15.09	−16.17	−11.85
2008		6.10	5.57	6.69	6.82	4.38	3.52
		15.98	14.03	11.40	−15.63	−10.83	−13.57
2009		9.42	7.30	5.82	5.71	6.58	5.57
		14.84	16.02	14.60	−14.69	−12.60	−9.49
2010		9.25	9.57	5.84	6.98	6.64	6.49
		17.13	13.42	10.88	−20.50	−12.07	−10.78
2011		9.98	8.23	7.18	5.20	5.38	7.18
		14.69	12.95	12.46	−16.12	−11.99	−10.08
2012		9.58	9.17	6.48	4.45	4.81	8.45
		16.13	15.43	13.40	−14.76	−11.46	−12.15
2013		11.84	10.13	7.84	4.19	4.82	6.58
		15.35	15.77	13.50	−13.60	−11.79	−11.20

Notes.- "Cohort" refers to the year of hiring. "Pr(transition)" shows the transition probabilities from part-time to full-time status and vice versa, relative to all job stayers between two periods. "Δhours" shows the mean change in hours worked upon transition between statuses.

TABLE C8: Gross transition rates of job stayers by tenure between full-time and part-time work, and associated mean change in hours worked

		PT to FT			FT to PT		
Cohort	Tenure:	1 year	2 years	3 years	1 year	2 years	3 years
2002-07	Pr(transition):	6.84	5.79	4.43	4.19	4.29	3.45
	Δhours:	16.46	14.47	12.73	−16.69	−14.57	−13.34
2008-13		9.31	8.22	6.64	5.66	5.56	6.03
		15.55	14.86	12.64	−15.48	−11.89	−11.20

Notes.- see Table C7.

TABLE C9: Share of mean hours changes among job stayers by tenure accounted for by transitions between full and part-time status and vice versa

	Tenure		
Cohort	1 year	2 years	3 years
2002-07	0.66	0.54	0.79
2008-13	0.60	0.68	0.61

Notes.- "Cohort" refers to the year of hiring. The shares not shown are accounted for by hours changes within full or part-time work.

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

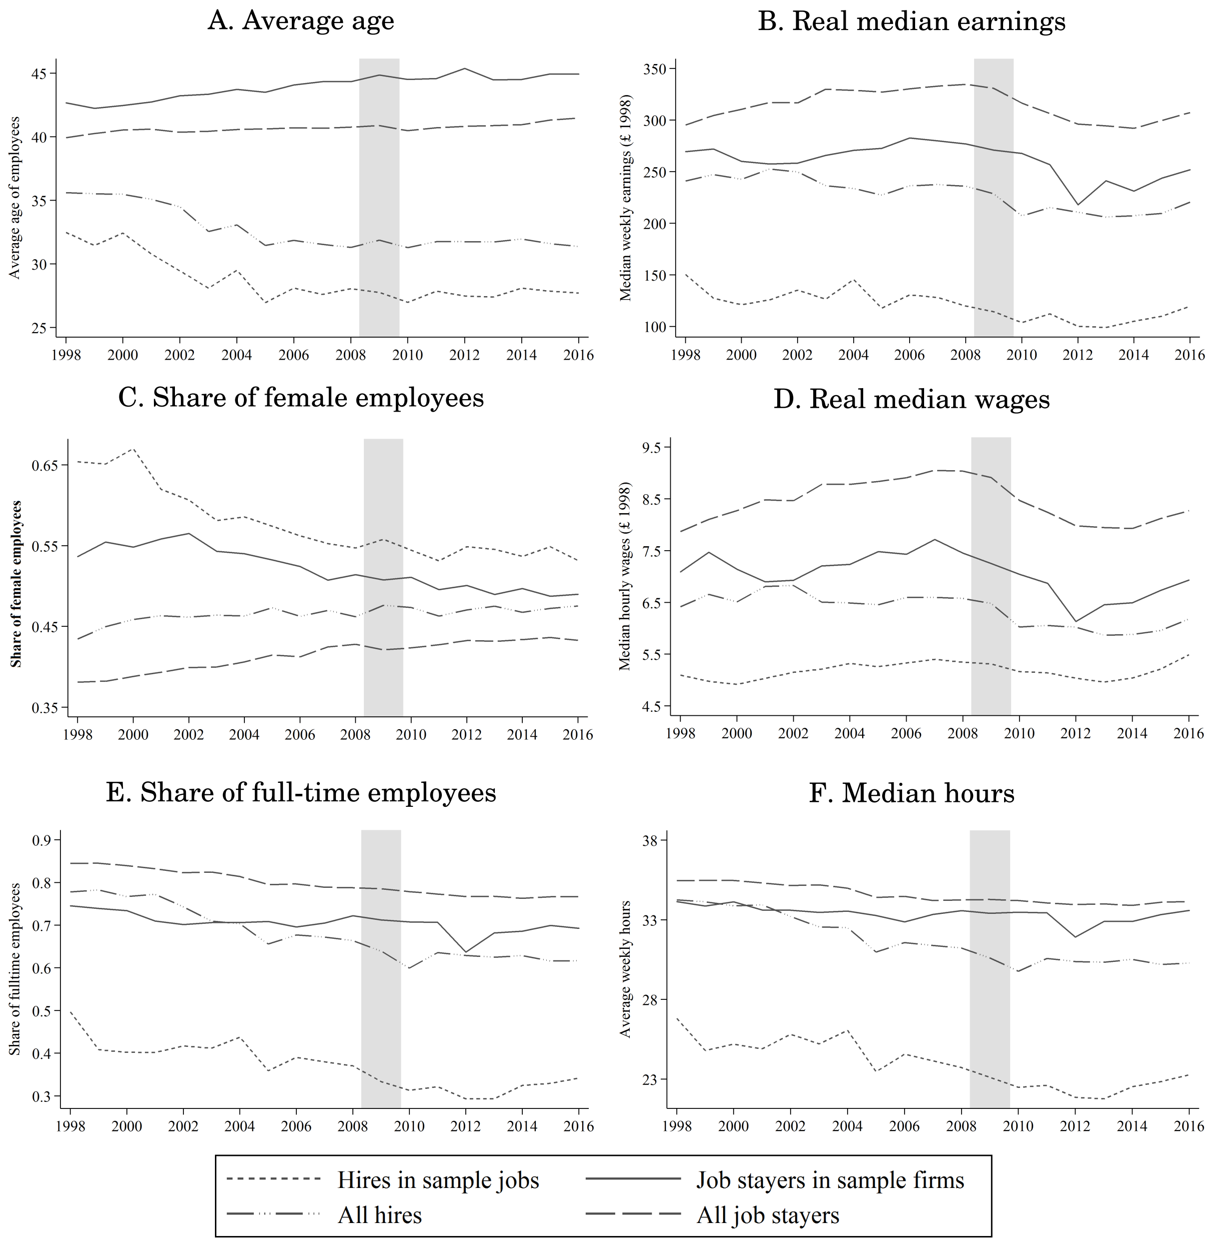

Notes.- shaded area marks official UK recession dates. “Hires in sample firms” refers to employees in entry-level jobs with less than twelve months tenure. “Job stayers in sample firms” are for jobs and employees who have more than 12 months tenure in the same job, and only for firms which are represented 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: Distribution of differences between log real wages of new hires and their median values within entry-level jobs, 1998-2016

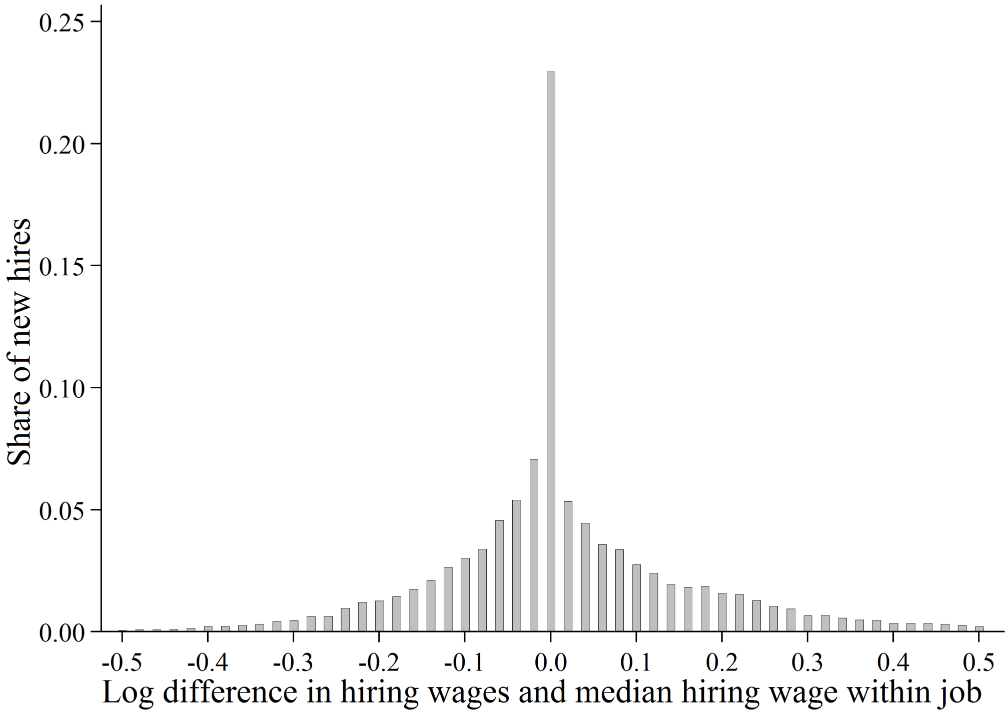

Notes.- within each entry-level job and year, the median hiring wage is subtracted from all hiring wages in that job and the resulting log differences are collected in bins with a width of two log points.

Figure D2 shows the distribution of employee hiring wages after subtracting their respective entry-level job median values. Some dispersion around the median hiring wages, indicated by zero, is visible here. We constructed similar figures for each year (omitted for brevity), and did not find evidence of systematic variation over the business cycle. In particular, the mass in the tails of the distribution does not change and the interquartile range is constant over the sample period.

FIGURE D3: Comparison of business cycle indicators, 1998-2016

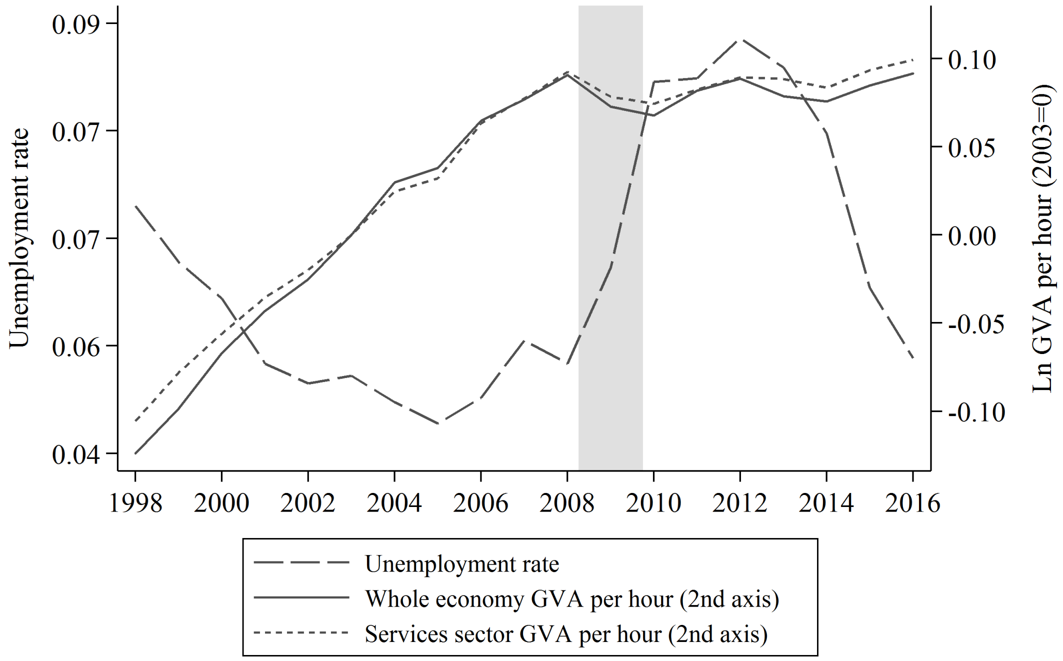

Notes.- see Section 2 and Table 6 for sources. Shaded area marks official UK recession dates.

FIGURE D4: 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

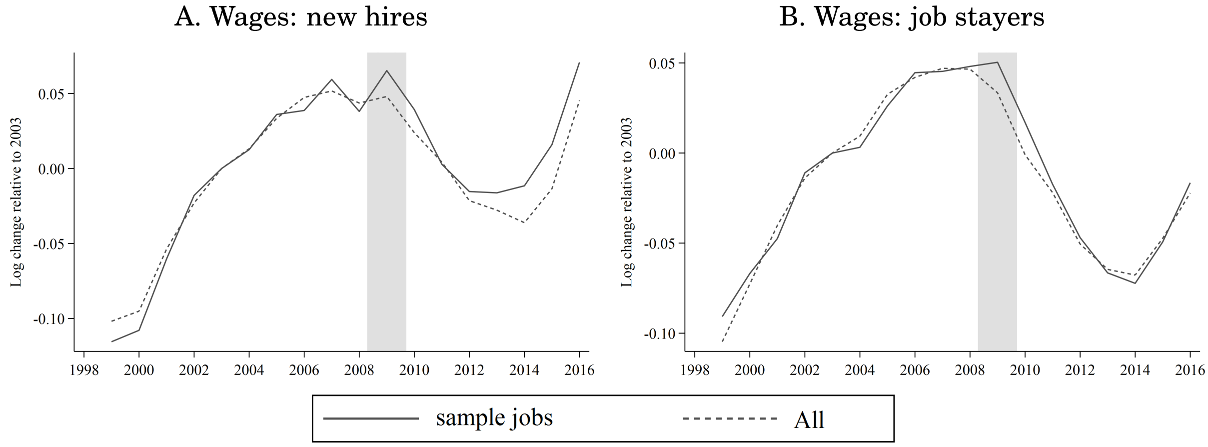

Notes.- see Section 2 for further details of sample construction. “All” here refers to all firms and jobs represented in the ASHE. Shaded area marks official UK recession dates.

FIGURE D5: Estimated period-fixed effects for hours worked, 1998-2016: comparison of entry-level jobs, all new hires, job stayers in the CH-firms sample, and all job-stayers

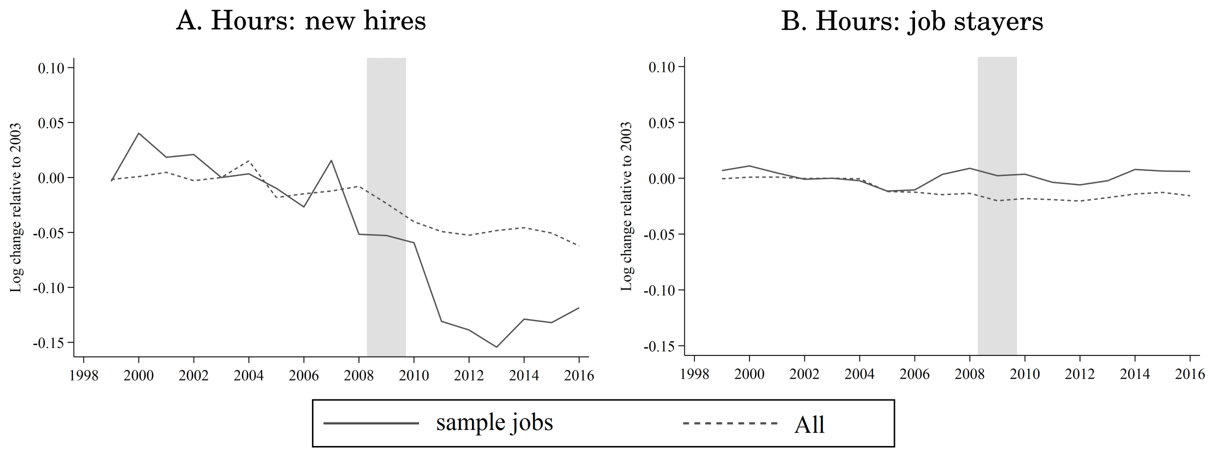

Notes.- see Section 2 for further details of sample construction. “All” here refers to all firms and jobs represented in the ASHE. Shaded area marks official UK recession dates.

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

The nominal wage changes of job stayers in the UK have been analyzed previously by Nickell and Quintini (2003) and most recently by Elsby et al. (2016). We briefly summarize results for year-to-year changes in the log nominal wages of job stayers in the baseline sample of consistent-hiring firms. Figure E1 shows the distributions of log changes for job stayers for the periods 2005-06 and 2009-10. These are representative of periods with relatively low (2005-06) and relatively high (2009-10) shares of job stayers with nominal wage cuts. Table E1 also displays summary statistics for all years in the sample. The dashed line marks the CPI rate of inflation 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 percent to 2.1 percent in the period before the Great Recession, as Table E1 shows. The distribution in Figure E1A 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 percent, appear to have experienced nominal wage cuts. This share increased during the recession to around 25 percent on average. Similarly, the share of job stayers with exactly zero nominal wage growth peaked at 5.7 percent between 2009-10. In particular, Figure E1B displays a relatively large share of nominal wage changes between zero and two percent for job stayers in CH-firms between 2009-10. However, the large increase in the share of job stayers which experienced negative changes in log *real* wages, as shown in the last column of Table E1, was mainly related to 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 percent of year-to-year job stayers appear to experience nominal wage cuts in the UK, suggesting that there is a relatively high degree of nominal wage flexibility in the British labor market. Nevertheless, the increase in inflation during the Great Recession resulted in over two-thirds of job stayers seeing their real wages cut.

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

A. 2005-06

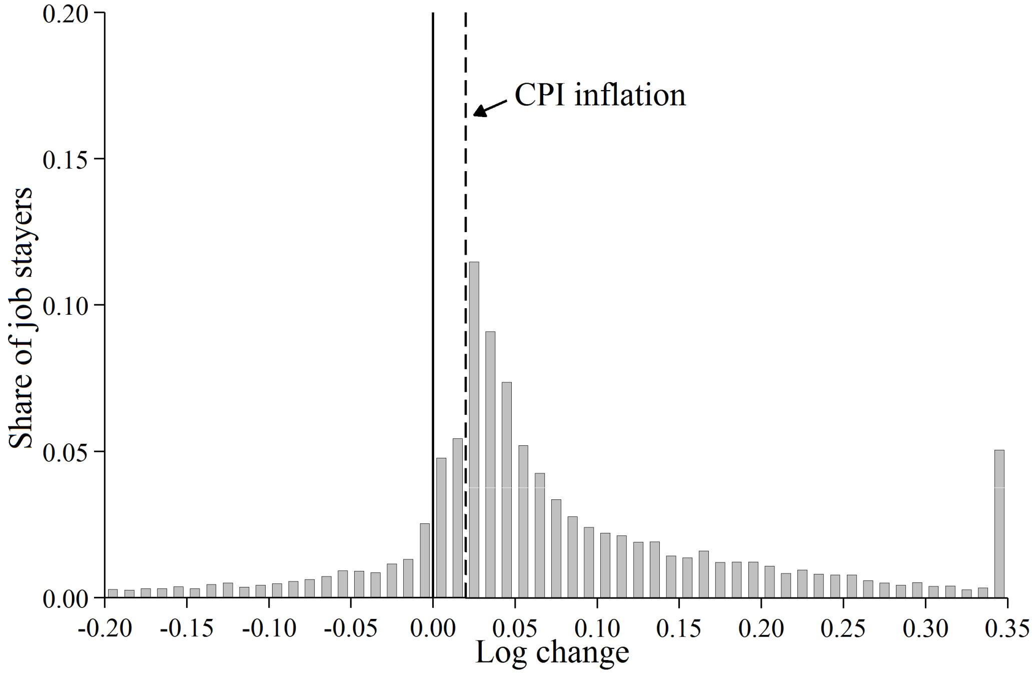

B. 2009-10

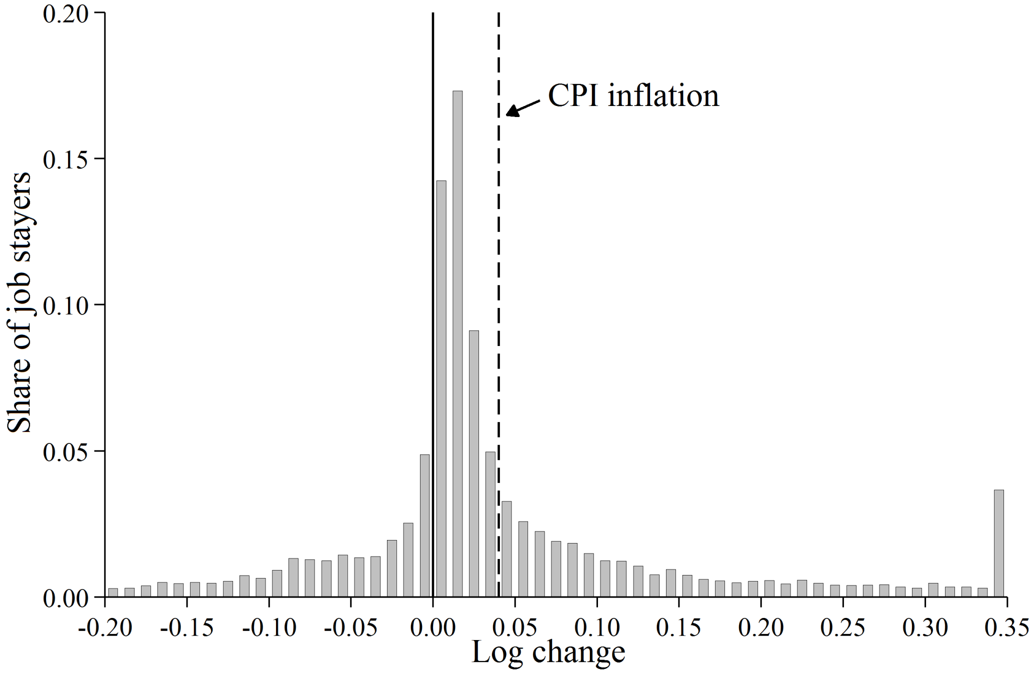

Notes.- solid line marks zero, dashed line marks the log change in the Consumer Price Index. Bars show half-open intervals, excluding the upper limit. Ages 16-64, private sector only.

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

	Percentage of log nominal wage changes 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 the average of log changes in CPI over the previous four quarters.

Appendix F. The role of the National Minimum Wage

Our results suggest that the real wages of new hires are equally responsive to business cycle conditions as job stayers' wages. 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 [Source: https://www.gov.uk/national-minimum-wage-rates; accessed 01/07/2017]. Collectively bargained wages can also limit a firm's flexibility in setting hiring wages. However, at the onset of the Great Recession, only six percent of new hires in the 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 F1: Real hourly wages of new hires and NMW adult rate, ages 22-64

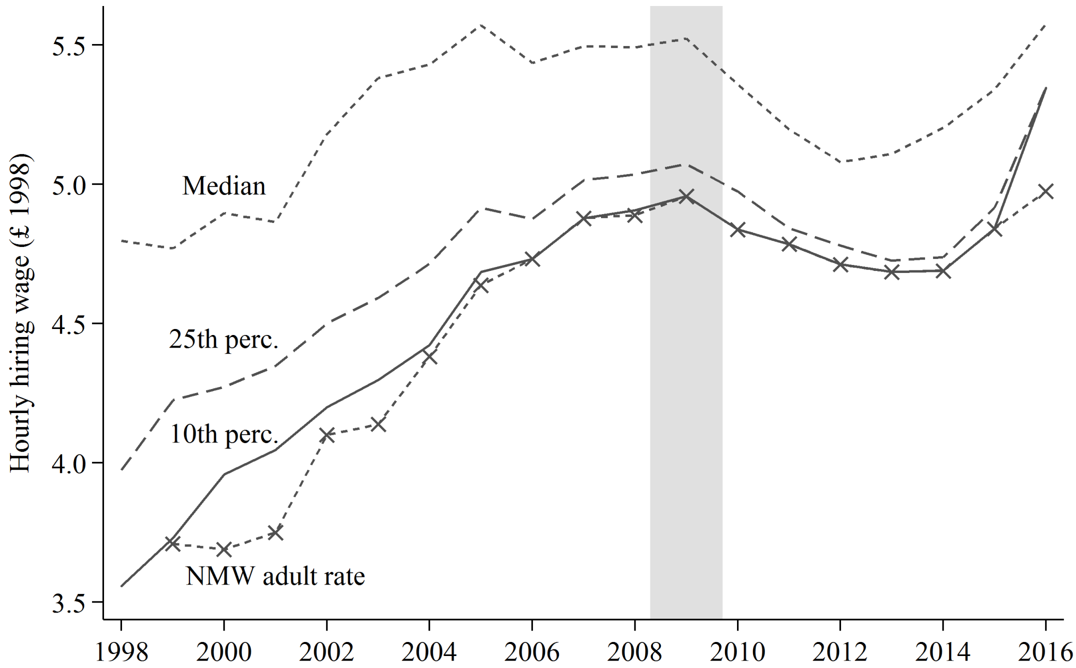

Notes.- National Minimum Wage adult rate and 10th, 25th, and median percentile of job-level hourly hiring wages, ages 22-64. All monetary values are deflated to 1998 values using the CPI. Shaded area marks official UK recession.

Figure F1 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 the baseline entry-level jobs sample for each year. The adult rate age limit was decreased from 22 to 21 in 2010. 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 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).

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 the next section below. In short, a partial equilibrium assumption underlies it: 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 F1), and the partial equilibrium assumption seems least restrictive. 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 the 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, 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 percent 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. For example, if under the counterfactual in 2013 ten percent of jobs hired at a wage below the NMW, then for the bottom decile of jobs ranked by actual hiring wages we would impute the hiring wages below the NMW from the counterfactual. Then we re-estimate regressions (1) and (2), using each of the counterfactual real hiring wage samples estimated with varying spillover parameter values. Figure F2 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 there is no spillover effect, the left-most circle shows that the responsiveness of real hiring wages to the unemployment rate increases from -2.56 to -2.64 percent. The standard errors are comparable to the baseline value (0.9) and lie outside the range of this figure. The semi-elasticity falls below -2.7 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 revert 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 firms in how far they could reduce the wages of new hires during the Great Recession.

FIGURE F2: Counterfactual estimates of the semi-elasticity of real hiring wages, 1998-2016: varying the assumed spillover effect of the NMW

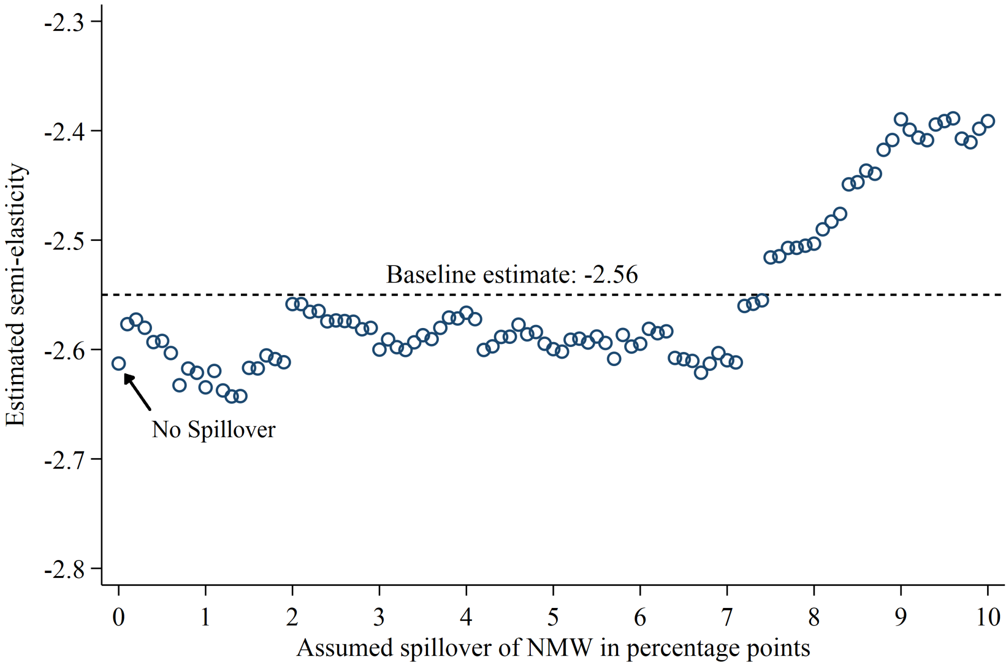

Notes.- each circle represents an estimate of the semi-elasticity of real hiring wages with respect to the unemployment rate. Standard errors lie outside of the figure. The horizontal axis shows the assumed spillover effect in p.p.. Dashed line shows the baseline estimate of semi-elasticity. We use a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

F.1 Description of the kernel re-weighting method

The following describes the method of DiNardo et al. (1996), which we use to estimate counterfactual densities of real hiring wages: the exposition here follows closely their own. 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 realized 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. The 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:*

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

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 F3, which shows estimated densities of real hiring wages (we pool data over multiple periods in Figures F3 and F4 for data confidentiality reasons). The kernel density estimate in the Figure F3A 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. Figure F3B shows the corresponding density for 2012-14, where the real NMW remained nearly constant around its 2013 value. The counterfactual density displayed in Figure F3C is a simple combination of the part of the density to the left of the solid line in Figure F3A and to the right of the solid line in Figure F3B, scaled to integrate to one.

Increasing the assumed spillover of the NMW acts as if shifting the NMW in Figure F3A-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 F4A 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 F4B displays the difference between the estimated and counterfactual wage densities shown in Figure F4A 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).

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 give

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

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 $\widehat{\psi}$ for each observation and use these weights in the kernel density estimation to derive the counterfactual density of real hiring wages in 2013. 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 $\widehat{\psi}_j$ 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.

FIGURE F3: Illustration of the re-weighting procedure for log real hourly hiring wages

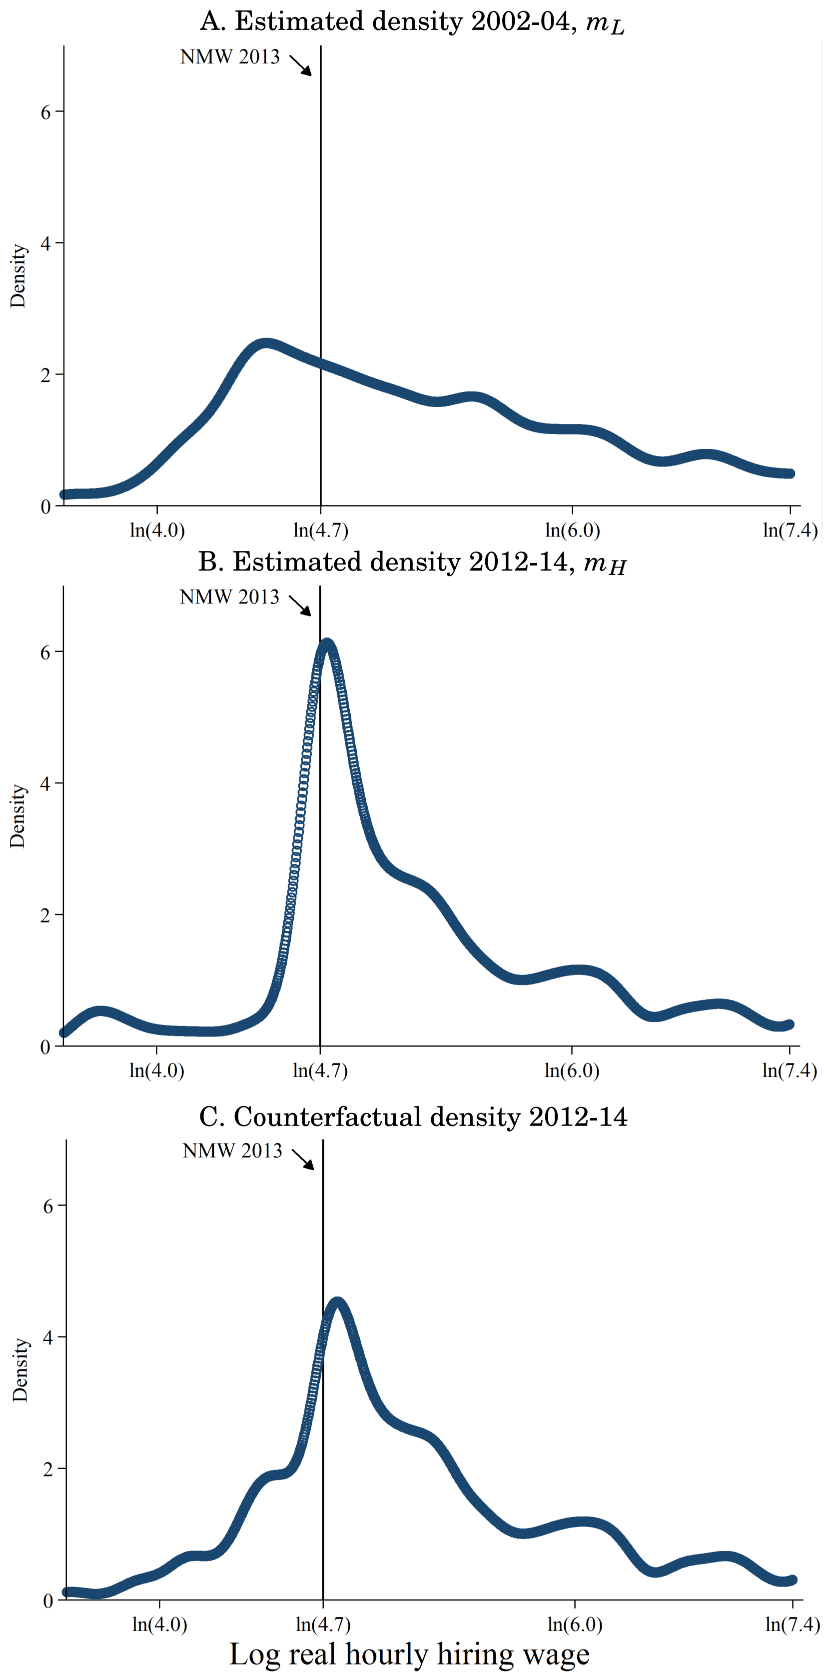

Notes.- densities estimated using Gaussian kernel and bandwidth of 0.03 (A) and 0.02 (B-C). Monetary values deflated to 1998 values using the CPI. Solid lines show the real value of the adult rate minimum wage in 2013.

FIGURE F4: Estimated and counterfactual densities of real log hiring wages at the job-level, 2012-14

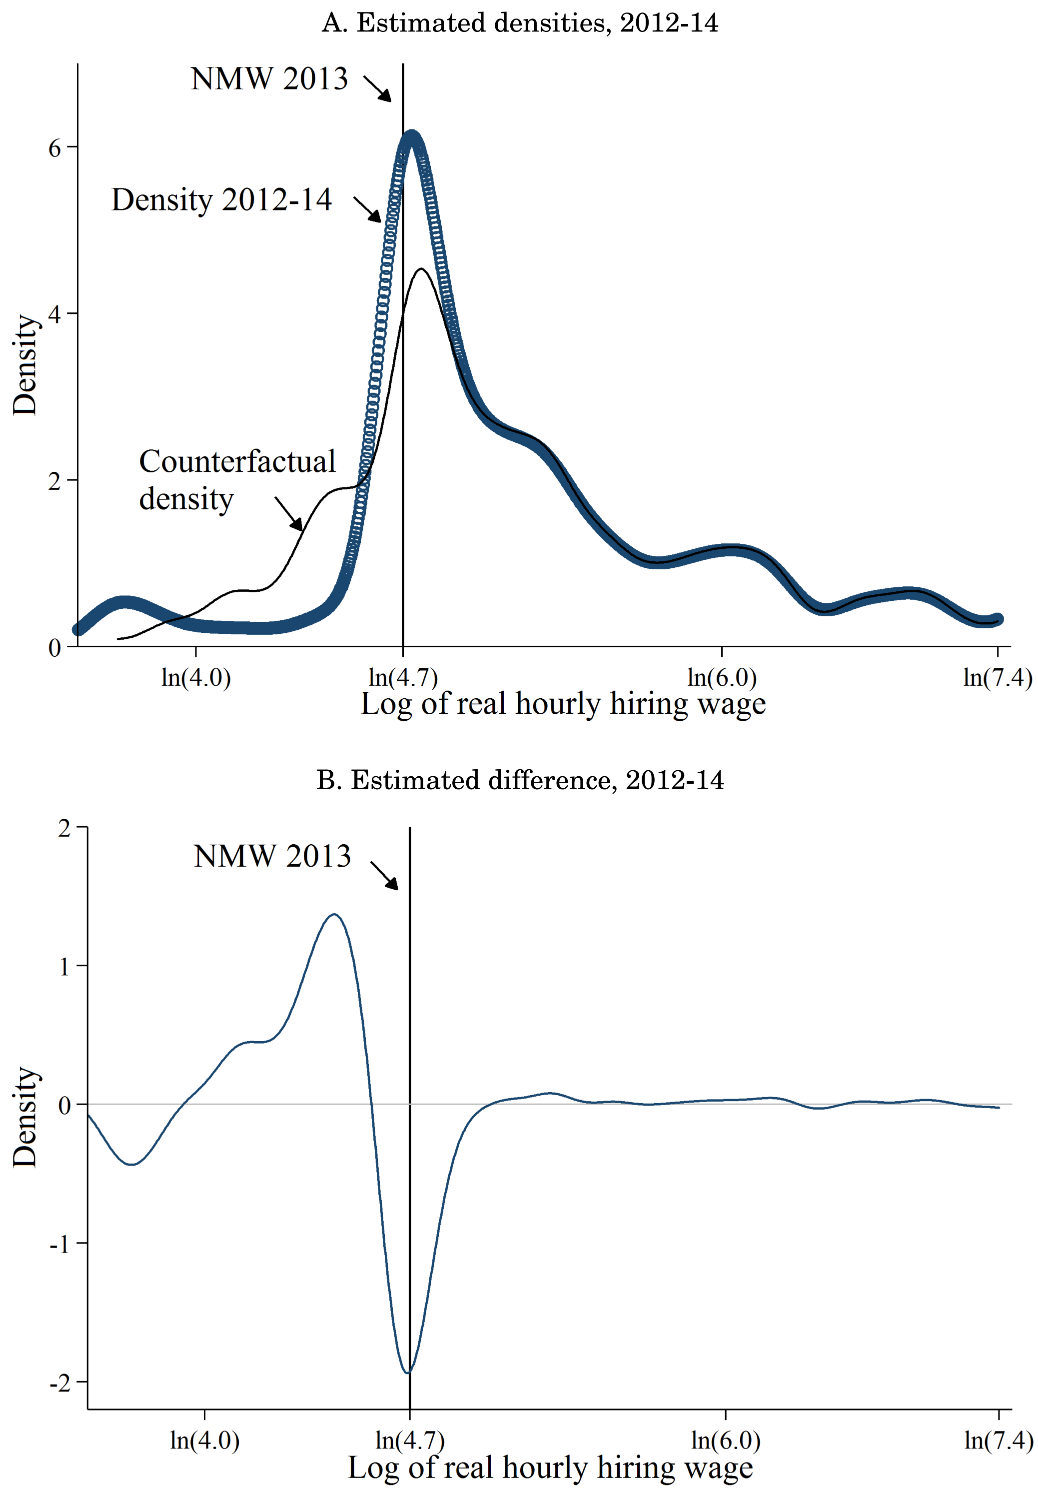

Notes.- densities estimated using Gaussian kernel and bandwidth of 0.02. Monetary values deflated to 1998 values using the CPI. Solid vertical line shows the real value of the adult-rate minimum wage.

Appendix G. Possible (UK-specific) 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) emphasize that the decline in unionization in the UK since the 1970s could only account for a small part of these observed changes in the behavior 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 unionization than the UK.

A similar line of argument applies for the role inflation. In both the US and the UK, price inflation was historically low before and during the Great Recession. In Online 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 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 percent of employees before 2008-9 to 3-5 percent in the years after. Our main findings on the extent of UK real wage flexibility reflect the fact that large numbers of employees appear to experience yearly nominal wage cuts, almost independently of the economic cycle.

Blundell et al. (2014) argue that the UK's labor 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 policy changes led to an increase of almost ten percent in the employment rate among UK lone parents, despite this occurring throughout a major recession (Avram et al., 2018). It is plausible that increased competition for jobs, brought on by the cumulative and extensive changes in the UK's active labor 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 behavior of the UK labor 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 labor 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 record: almost 1.5 million (6 percent of all employees), compared with 2.5 million unemployed, and compared with 0.7 million involuntary part-time employed in 2007 [Source: ONS Labour Market Statistics, October 2017, available at https://www.ons.gov.uk/; accessed 07/11/2017. See also Bell and Blanchflower (2013) for more details about so-called "Underemployment" in the UK].

Another possible cyclical feature of labor markets is the so-called "Added Worker Effect", whereby individual household members will increase their labor 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 labor force transition rate data (Razzu and Singleton, 2016). However, Bryan and Longhi (2018) 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 behavior in the presence of fixed costs among US establishments. Moreover, if firms have to pay 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.

Online Appendix References

Avram, S., M. Brewer, and A. Salvatori. 2018. "Can't work or won't work: Quasi-experimental evidence on work search requirements for single parents." *Labour Economics*, 51 63 – 85.

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

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

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

Bryan, M., and S. Longhi. 2018. "Couples' Labour Supply Responses to Job Loss: Growth and Recession Compared." *Manchester School*, 86(3): 333–357.

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.

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." *The Economic Journal*, 114(494): 95–101.

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.

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.

Gregg, P., S. Machin, and M. Fernández-Salgado. 2014. "Real Wages and Unemployment in the Big Squeeze." *The Economic Journal*, 124(576): 408–432.

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

Nickell, S., and G. Quintini. 2003. "Nominal Wage Rigidity and the Rate of Inflation." *The Economic Journal*, 113(490): 762–781.

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.

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.

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.