Department of Economics

The Extent of Downward Nominal Wage Rigidity: New Evidence from Payroll Data

By Daniel Schaefer and Carl Singleton

Discussion Paper No. 2021-22

Department of Economics
University of Reading
Whiteknights
Reading
RG6 6EL
United Kingdom

https://www.reading.ac.uk/economics/

The Extent of Downward Nominal Wage Rigidity: New Evidence from Payroll Data

Daniel Schaefer

Johannes Kepler University Linz

Carl Singleton*

University of Reading

30th November 2021

Abstract

Low inflation has forced the topic of downward nominal wage rigidity (DNWR) back to the centre stage of macroeconomics. We use over a decade of representative payroll data from Great Britain to document novel facts about wage adjustments. We find that basic wages drive the cyclicality of marginal labour costs, which makes them the most relevant wage measure for macroeconomic models that incorporate wage rigidity. Basic wages show substantially more evidence of downward rigidity than previously documented. Every fifth hourly-paid and every sixth salaried employee normally sees no basic wage change from year-to-year, and very few experience cuts. Wage freezes were more common in the Great Recession and are far more likely in smaller firms. We also find evidence that employers compress wage growth when inflation is low, indicating that DNWR constrains wage setting. Further, we show that the wages of new hires and incumbent employees respond equally to the business cycle. These results all point to the importance of including DNWR in macroeconomic and monetary policy models, and our simulations demonstrate that the empirical extent of DNWR can cause considerable long-run output losses.

Keywords: Downward nominal wage rigidity; Hiring wages; Unemployment fluctuations; Macroeconomic policy; Marginal labour costs
JEL codes: E24, E32, J31, J33

*Schaefer: Department of Economics, Johannes Kepler University Linz, Altenberger Straße 69, 4040 Linz, Austria (email: daniel.schaefer@jku.at; corresponding author). Singleton: Department of Economics, University of Reading, Whiteknights, RG6 6EL, UK (email: c.a.singleton@reading.ac.uk). This paper builds on our previous working paper in Schaefer and Singleton (2020*a*).

We are grateful for helpful comments and advice from Mark Casson, Mike Elsby, Marco Fongoni, Jan Grobovšek, Jochen Güntner, Robert Hart, Jonathon Hazell, Chris Martin, Tho Pham, and Gary Solon, as well as participants in 2021 at the Annual Meeting of the Austrian Economic Association (Innsbruck) and the European Association of Labour Economists Annual Meeting (Padua).

This work is based on the Annual Survey of Hours and Earnings Dataset (Crown copyright 2019), 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.

Singleton thanks the Economic and Social Research Council and Administrative Data Research UK for the funding support of the Wage and Employment Dynamics project (ES/T013877/1).

1. Introduction

Over eighty years have passed since Keynes (1936) proposed that workers simply refuse to accept nominal wage cuts, thus preventing the real price of labour from falling during recessions and leading to rising unemployment. An implication of Keynes' theory is that moderate inflation can *"grease the wheels of the labor market"* (Tobin, 1972) by bringing down real wages and boosting employment. With stubbornly low inflation having become a persistent feature in many developed economies, the question whether labour markets display signs of downward nominal wage rigidity (DNWR) has attracted renewed attention. The answer has wide-ranging consequences for policy-makers; macroeconomic interventions can be justified if excess supply in the labour market does not self-correct through wage adjustments (Schmitt-Grohé and Uribe, 2016). Furthermore, the effects of monetary policy on the real economy depend crucially on the degree of nominal wage rigidity in the workhorse model of monetary policy analysis (Christiano, Eichenbaum, and Evans, 2005).

Despite being one of the central issues in macroeconomics, the empirical extent of nominal wage rigidity remains an open question. The key reason why economists disagree is that studies on DNWR have used datasets that were ill-suited to the task. The findings from household survey data have frequently been discounted on the grounds that self-reported wages contain substantial response errors, which can bias the results.[1] Recent studies have turned to more accurate administrative or payroll data, which suggest that wages are far more flexible than previously thought. The wealth of evidence has indeed become so great that researchers now question the often invoked assumption of DNWR (e.g., Elsby and Solon, 2019). But the datasets that have been used in these recent studies typically give total earnings, which consist of basic wages and extra payments, such as overtime or commission. This makes it tricky to interpret any results and inform macroeconomic models; nominal wage rigidity is typically incorporated to dampen the cyclical response of marginal labour costs, hence only the pay components that drive the cyclical responsiveness of the marginal costs actually matter for these models.[2]

Our study overcomes these challenges and provides novel empirical evidence on the extent of DNWR, using a unique longitudinal dataset from Great Britain, the Annual Survey of Hours and Earnings (ASHE), which offers six main advantages. First, the dataset comprises a one percent random sample of income tax-paying workers, allowing us to derive results that are representative of the entire labour market. This matters because we will document

[1] Studies on nominal wage rigidity using US household surveys include: McLaughlin (1994); Akerlof, Dickens, and Perry (1996); Card and Hyslop (1996); Kahn (1997); Altonji and Devereux (2000); Lebow, Saks, and Wilson (2003), and more recently Barattieri, Basu, and Gottschalk (2014); Elsby, Shin, and Solon (2016). Smith (2000) and Fehr and Goette (2005) similarly analysed household survey data from Great Britain and Switzerland, respectively. Dickens et al. (2007) provide results for a sample of European countries and the US.

[2] Administrative data on total earnings show little evidence of DNWR, except in countries where nominal wage cuts are legally impossible, e.g., Portugal and Sweden (see the recent survey by Elsby and Solon, 2019).

substantial cross-sectional heterogeneity over workers and firms in the extent of DNWR; for example, smaller firms show more signs of DNWR than larger firms. Second, this survey is administered to employers who are legally obliged to report information from their payroll, making the data more accurate than those obtained from household surveys (Elsby, Shin, and Solon, 2016). Third, employers are explicitly asked whether an employee has been working in the same job and role for more than a year, meaning that we can study actual *job* stayers instead of firm stayers. This distinction matters because we are interested in a firm's tendency to reduce its employees' pay for the same work. Fourth, the data provide detailed payroll-based records on basic wages, hours worked, and extra pay components. These allow us to study basic wages *per hour* separately from extra pay components, comparing like-for-like measures of hourly pay over time, as well as to analyse separately the cyclicality of each pay component. Fifth, we can distinguish between hourly-paid and salaried (non-hourly-paid) employees. Previous research has found conflicting evidence on the extent to which DNWR appears to differ between salaried and hourly-paid workers, questioning whether studies based on only hourly-paid workers generalise to the aggregate labour market.[3] Sixth and finally, the data cover a period of substantial macroeconomic instability, with significant variation in unemployment and low inflation. This matters, since there are reasons to expect that high inflation would hide clear signs of DNWR over the business cycle (Card and Hyslop, 1996).

Our study makes four main contributions. First, we confirm recent findings from US payroll data by Grigsby, Hurst, and Yildirmaz (2021) that basic wages are highly procyclical, whereas all other extra pay components are approximately constant over the business cycle. We also show that changes in an employee's basic wages are longer-lasting than changes in extra pay, implying that the former have the greatest impact on labour costs. Combined, these findings demonstrate that the measure of basic wages from payroll data is the relevant empirical counterpart to the notion of a wage in macroeconomic models that incorporate nominal wage rigidities to generate cyclical fluctuations in unemployment.

Second, we present evidence that basic wages are rigid downwards. Among employees who are constantly employed in the same job from year-to-year, the distributions of wage changes are markedly more asymmetric than previously thought in Great Britain, falling off sharply to the left of zero. Additionally, we combine the theoretical insights of Elsby (2009) and our payroll data to answer whether DNWR was constraining firms' wage setting. His theory predicts that forward-looking firms, if constrained by DNWR, will refrain from raising nominal wages today, because this increases the likelihood of having to cut wages, at some cost, in the future. We use unconditional quantile regression to show that job-stayer wage growth is indeed considerably compressed relative to a world without DNWR. Taken together,

[3] For example, Card and Hyslop (1996) analysed the Current Population Survey, concluding that DNWR is not lower for salaried than for hourly-paid workers in the US. Contrary to this, Kahn (1997) found that hourly wage rates exhibit substantially more signs of DNWR than salaries in the US Panel Study of Income Dynamics.

our results support the degree of nominal rigidity typically assumed in New Keynesian models (e.g., Christiano, Eichenbaum, and Evans, 2005), as well as the assumption of DNWR invoked in recent macroeconomic models of business cycle fluctuations (e.g., Daly and Hobijn, 2014; Dupraz, Nakamura, and Steinsson, 2019). Our findings also imply that DNWR causes sizeable output losses. To demonstrate this, we calibrate the dynamic stochastic general equilibrium model with downward rigid nominal wages of Benigno and Ricci (2011) to match our empirical estimates on the prevalence of wage freezes and cuts. Simulations of this model suggest that the *long-run* output loss caused by DNWR is around 0.7 to 1.3 percent of GDP in a low-inflation environment. These substantial economic costs suggest that the optimal rate of inflation may be significantly positive rather than zero or negative (the Friedman rule). For example, increasing the equilibrium rate of inflation from 1 to 4 percent per year would reduce the long-run output loss due to DNWR from 1.3 to 0.2 percent of GDP.

According to the literature on labour markets with search frictions, the key determinant of a firm's hiring decisions, and consequently the determinant of aggregate unemployment fluctuations, is the flexibility of the wages of new hires (Pissarides, 2009). Our third contribution shows that the wages of new hires are just as cyclical as the wages of incumbent employees, indicating that firms are constrained by internal wage structures. This implies that the flexibility of incumbent wages is a sufficient statistic for the flexibility of the wages of new hires, and that it is reasonable for macroeconomists to use DNWR to account for observed fluctuations in unemployment.

Our fourth contribution documents substantial heterogeneity in the extent of DNWR within payroll data. We find that firm size is especially significant in accounting for the extent of DNWR. In small firms with up to 50 employees, the conditional likelihood of a year-to-year wage freeze for job stayers is 27 percent for salaried workers and 34 percent for hourly-paid workers. These values drop to 15 percent in firms with more than 5,000 employees. These differences between jobs, as well as others discussed later, such as across industry sectors or depending on whether wages are affected by collective bargaining, highlight the importance of using nationally representative data when assessing the prevalence and macroeconomic significance of DNWR.

There are three previous studies of DNWR in Great Britain, which are highly relevant to our own. Smith (2000) found that cuts in weekly earnings were quite common among a subsample of respondents to the 1991-96 waves of the British Household Panel Survey who were prompted to check their payslips when asked about pay. Nickell and Quintini (2003) studied the precursor to the datasets that we are using, the New Earnings Survey for 1975-99, which neither identified hourly-paid workers, nor new hires, nor separated basic wages from extra pay components besides overtime. Nickell and Quintini confirmed Smith's findings: job-stayer earnings were frequently cut in Great Britain in the early 1990s. Recent work by Elsby, Shin, and Solon (2016) updated the study of Nickell and Quintini, finding that job-stayer earnings per hour, excluding overtime, are frequently cut in Great Britain and earnings freezes do not occur excessively often, concluding that DNWR may be less binding than is often supposed. We replicate those findings, but we also show that the ASHE survey design likely leads to previously undocumented recording errors in the data on hours worked, implying that earlier studies may have under-estimated the degree of DNWR.

Three research teams have recently investigated the extent of DNWR in US administrative or payroll data for job stayers.[4] Jardim, Solon, and Vigdor (2019) and Kurmann and McEntarfer (2019) used administrative data from the State of Washington. Unusually for the US, Washington requires firms to report hours worked, and so these authors could derive measures of total earnings per hour. Both studies have documented a considerable proportion of cuts among job-stayer total earnings per hour, and Jardim, Solon, and Vigdor also found this in a subsample of employees where overtime was likely to have been rare. Their results provide important insights for total labour costs, but they cannot separate cyclical from non-cyclical pay components, nor do they identify hourly-paid workers or new hires. Using a proprietary dataset from a US payroll processing firm, Grigsby, Hurst, and Yildirmaz (2021) found significantly more evidence of nominal rigidity in basic wages than was previously documented. They also confirmed that the wages of new hires and incumbents are equally cyclical. We extend their insightful work in two important ways: first, we analyse whether nominal wages are rigid downwards and assess what the consequences of DNWR for the aggregate economy might be, and second, we study basic wages for a sample of jobs that is representative of an entire national labour market. Importantly, as Grigsby, Hurst, and Yildirmaz acknowledged, the data they used under-represented very large firms. We show that this could partially explain the relatively high share of basic wage freezes observed in their data compared with what we find in Great Britain. Nevertheless, we can confirm many of their results, suggesting that their findings are not driven by idiosyncrasies of the US labour market. Thus, we provide valuable additional evidence on downward rigidity of nominal basic wages.[5]

The rest of the paper is structured as follows: Section 2 describes the data and our sample selection; Section 3 argues that basic wages drive (marginal) labour costs; Section 4 analyses job-stayer basic wages, separately for hourly-paid and salaried workers, and documents wage growth compression; Section 5 uses regression analysis to show that the responsiveness of job-stayer and new-hire wages to the business cycle are not significantly different; Section 6

[4] Hazell and Taska (2020) find substantial evidence of DNWR for new hires within online vacancy data from the US. They argue that this evidence is consistent with DNWR at the job level being the key for unemployment fluctuations. They do not study wages of incumbent employees.

[5] Carneiro, Portugal, and Varejao (2014) studied Portuguese administrative microdata and found strong signs of DNWR, with around 45 percent of job stayers having exactly the same basic wage in 2012 as in 2011. However, as the authors explained, nominal wage cuts are explicitly prohibited by law in Portugal. Sigurdsson and Sigurdardottir (2016) analysed Icelandic payroll-based data, where they observed that wage cuts were rare and the share of year-to-year basic wage freezes for job stayers was around 16 percent. These studies neither demonstrated that hiring wages were as flexible as job-stayer wages, nor were they analysing whether firms were constrained by DNWR and the consequences for the aggregate economy.

uncovers heterogeneity across workers and firms in the conditional probability of basic wage freezes and cuts; Section 7 discusses the macroeconomic implications; and Section 8 concludes.

2. Description of the Annual Survey of Hours and Earnings

Our analysis of DNWR uses the Annual Survey of Hours and Earnings (ASHE) (Office for National Statistics, 2019). The ASHE is an ongoing longitudinal panel of employees, starting in 2004, based on a one percent random sample of workers in Great Britain who pay income tax or make National Insurance contributions. The last observations in our study are from 2018. Employers respond to the survey by providing information from the pay period that includes a specific date in April, either by returning a survey questionnaire or directly through their payroll by a special arrangement with the Office for National Statistics (ONS). This setup implies that we only have data each year for individuals in the panel who were employees on the survey reference dates.[6]

The ASHE offers a unique combination of features which make it ideal for this study. First, employers are legally obliged to report employee earnings with reference to payrolls, making the data more accurate than those obtained from household surveys (Nickell and Quintini, 2003; Elsby, Shin, and Solon, 2016). Second, employers are explicitly asked whether an employee has been working in the same job and role for more than a year, meaning that we can study actual *job* stayers instead of firm stayers. This distinction matters, because we are interested in a firm's tendency to reduce its employees' pay for the same work; a promotion or change in the job may cause some adjustment of an employee's wage, but there is little reason to expect much nominal rigidity for such job changers. Third, the ASHE has a large sample size, with up to 100,000 wage-change observations of job stayers per year. Finally, the ASHE is representative of jobs in Great Britain, being based on a one percent random sample of employees.[7]

Relative to its precursor, the New Earnings Survey Panel Dataset (NESPD), the ASHE contains three major improvements. First, it identifies workers who received an hourly pay rate, hereafter 'hourly-paid' workers. We refer to all other employees as 'salaried'. This allows us to assess how far results based on hourly-paid workers generalise to the aggregate labour market, as is often implicitly assumed to be the case (e.g., Barattieri, Basu, and Gottschalk, 2014). Second, the ASHE contains details on a worker's composition of pay,

[6] We only use data starting from 2006, because questionnaire changes in 2005 and 2006 introduced inconsistencies between these years. Specifically, new instructions were included on how firms should report employee hours worked. See Schaefer and Singleton (2019, 2020*b*) for further descriptions of the dataset.

[7] The ASHE does not include the very low-paid employees; around 3-4 percent of employees in Great Britain (ONS estimates of this undercoverage). The design of the sampling frame implies that it contains no observations of employees with earnings below the National Insurance threshold or who work for businesses that have a turnover below the Value Added Tax (VAT) threshold, e.g., £5,564 and £73,000 per year in 2012-13, respectively.

reporting separately basic pay and extra pay components. The latter include overtime pay, shift premium pay, incentive pay received for work carried out in the April pay period, and other pay (e.g., meal or travel allowances).[8] We use the information on basic pay and hours worked to compute salaried workers' basic wages. For hourly-paid workers, the basic wage is just their reported hourly basic pay rate. We provide more details on the pay components and their exact definitions in Appendix A. Third, the ASHE reports when an employee starting working for a firm, which allows us to distinguish new hires from incumbent employees.

We study employees aged 16-64, who did not incur any loss of pay in the April reference period (e.g., unpaid sick leave) and who were not paid at an apprenticeship or a trainee rate. We drop person-year observations if a worker held multiple jobs, was reported as having worked on average less than one or more than 100 hours per week in April, was reported as being paid less than 80 percent of the age-relevant statutory National Minimum Wage, or had missing or imputed values for any of the pay variables which we are interested in. Altogether, our sample selection criteria result in a working dataset of 1,843,172 employee-year observations over our 13-year sample period, 2006-2018. Appendix A describes and justifies our sample construction in more detail. We define a 'job stayer' as an employee whom we observe working in the same job as in the previous April, such that we can measure within-job year-to-year wage changes.

Table 1 shows descriptive statistics for the samples of job stayers and all employees in the working dataset, where the latter includes job switchers, (re-)entrants to employment and job stayers in 2006. On average, the ASHE suggests that around 92 percent of British employees who remain in employment from one year to the next are job stayers, as opposed to job switchers.[9] Hourly-paid job stayers are more likely to be working for a private sector company, are slightly younger, and are less likely to be employed on a full-time basis than salaried job stayers. Around half of all job stayers are covered by a collective agreement, defined as any arrangement affecting the pay of more than one employee. The median and average basic wages are higher for salaried job stayers. Firms in the industry sectors of wholesale & retail trade and hotels & restaurants represent 32 percent of hourly-paid job stayers, while only 13 percent of salaried job stayers work in those industries. Professional services and financial services employ around 42 percent of the salaried job stayers (see Appendix Table A2).

Basic wages are the primary source of labour income for the vast majority of job stayers: on average, 94 percent of an employee's income is accounted for by basic earnings (basic wages times basic hours worked), and over half of all job stayers have no other labour income besides basic wages. Figure 1 shows the share of basic earnings in total earnings (basic plus

[8]Using the NESPD, the previous studies by Nickell and Quintini (2003) and Elsby, Shin, and Solon (2016) were restricted to analysing a worker's total earnings per hour, excluding overtime. Additionally, these authors could neither distinguish hourly-paid from salaried workers, nor identify new hires.

[9]We can observe a small minority of job stayers who switch between being hourly-paid and salaried, but all our analysis only considers job stayers who are paid in the same way in consecutive years. This does not affect any of our results.

TABLE 1: Descriptive statistics of employees in Great Britain, Annual Survey of Hours and Earnings, 2006-18

	Job stayers		All employees
	Hourly-paid (I)	Salaried (II)	(III)
Private sector (%)	76.3	60.5	66.5
Female (%)	51.8	49.8	51.1
Age (years)	42.6	41.7	40.7
Full-time (≥ 30 hours, %)	65.4	79.7	73.0
Collective agreement (%)	49.3	49.5	46.6
Firm size (no. of employees, median)	2,570	2,640	2,200
Weekly basic hours (median)	36.8	37.0	37.0
Basic wage (mean, £)	10.65	15.42	13.80
Basic wage (median, £)	8.91	12.95	11.10
N	320,087	594,709	1,843,172

Notes: Basic wages are deflated to 2015 GB Pounds (GBP) using the UK Consumer Price Index. Firm size rounded to nearest ten for statistical disclosure control.

extra earnings) along the basic wage distribution. To generate this figure, first we group job stayers into percentiles of the basic wage distribution, and then compute the average share of basic earnings in total earnings within each percentile. The share of basic earnings in total earnings is increasing with the level of the basic wage, reaching over 95 percent in the top percentile for both hourly-paid and salaried job stayers.

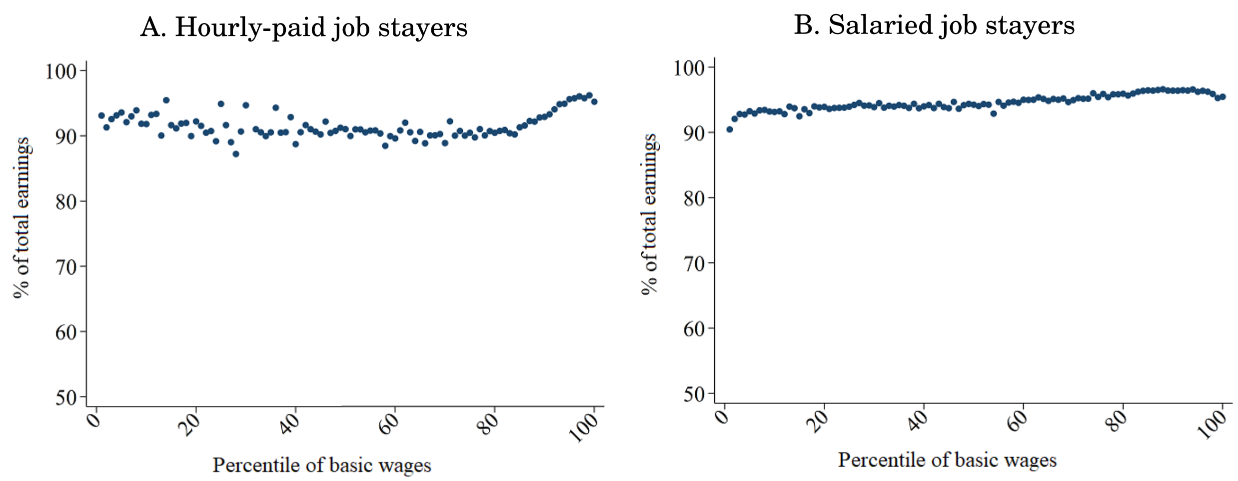

FIGURE 1: Importance of basic earnings along the basic wage distribution

Notes: Average shares of basic earnings in total earnings within the corresponding percentile of the basic wage distribution. Data are pooled across the sample period 2006-2018.

3. The Cyclicality of Basic Wages and Extra Pay Components

Nominal wage rigidity is typically incorporated into macroeconomic frameworks to dampen the cyclical movements of firms' marginal cost of labour (e.g., Christiano, Eichenbaum, and Evans, 2005; Smets and Wouters, 2007). Intuitively, if rigidities prevent the nominal marginal cost of labour from falling sufficiently in response to a negative shock, then firms will produce less output and demand less labour. Therefore, the relevant wage concept for macroeconomic models is best captured by the pay components with the largest impact on the cyclicality of marginal labour costs. In the previous section, we showed that basic wages are the only income source for the majority of workers and make up 94 percent of all labour income, on average. However, basic wages would only be the most relevant variable for the cyclicality of marginal labour costs if they were procyclical, while the extra pay components did not move systematically with the business cycle (Grigsby, Hurst, and Yildirmaz, 2021).

To understand which pay components are the most cyclical, we estimate their response to regional unemployment rates, a proxy for the state of the business cycle, using least squares:

$$\Delta \log(w_{ijrt}) = \theta_{ij} + \beta^u U_{rt} + \mathbf{x}'_{it}\boldsymbol{\delta} + \varepsilon_{ijrt} \ , \tag{1}$$

whereby w_{ijrt} are the various potential measures of wages for individual i, who works in job j, in region r, and in year t. The unemployment rate (in percent) in region r and year t is U_{rt}, and θ_{ij} is a fixed effect for a match between employer and employee. The coefficient of interest is β^u, which measures the semi-elasticity of wage growth to the unemployment rate. The vector $\mathbf{x}_{it}$ contains time-varying controls for firm size, employee age and its square, and tenure squared.[10] The regions are the eleven EU-NUTS1 administrative regions of Great Britain (e.g., London, Wales, Scotland, North West).

Table 2 displays the estimates of β^u in equation (1), separately for salaried and hourly-paid employees. Our results indicate that basic wages are considerably procyclical. When the regional unemployment rate increases by one percentage point, nominal basic wage growth is on average 0.46 percent lower for salaried workers and 0.38 percent lower for hourly-paid workers. Given the UK unemployment rate increased by over 3 percentage points during the Great Recession, the estimates of β^u imply that expected wage growth was lower by 1.38 and 1.14 percent for salaried and hourly-paid job stayers, respectively. These effects are quantitatively large, given that average annual nominal wage growth was only 3.8 percent and 3.5 percent for job stayers over the sample period.

Column (II) of Table 2 shows the estimates when adding the extra pay components, except overtime pay, to basic wages. The slightly smaller point estimates are neither statistically nor economically different from those focusing only on basic wages. In the last column of

[10] A linear control for tenure is excluded, because it would be perfectly collinear with an employee's age within an employer-employee match.

TABLE 2: Estimated cyclicality of basic wages and extra pay components for job stayers

	Basic wages (I)	Basic wages plus shift, incentive, and other pay (II)	As (II) plus overtime (III)
Salaried	−0.456	−0.476	−0.498
	(0.032)	(0.034)	(0.036)
Hourly-paid	−0.379	−0.413	−0.413
	(0.037)	(0.034)	(0.049)

Notes: Least squares estimates of the semi-elasticity of wages with respect to the regional unemployment rate (in %), β^u, in Equation (1), separately for salaried and hourly-paid job stayers. Controls are firm size, employee age, age squared, and tenure squared.
Standard errors in parentheses robust to three-way clustering over year, NUTS1 region, and match.
Sample sizes: Salaried job stayers, 605,680, hourly-paid job stayers, 327,109, after dropping singletons.
Sources: NUTS1 regional unemployment rates are from the Office for National Statistics (ONS) for April of each year, corresponding to the reporting period of the ASHE.

Table 2, we add overtime pay, such that the coefficient estimates show the cyclical responses of total earnings per hour. These estimates are 0.03-0.04 percentage points smaller compared to those in column (I), again neither statistically nor economically significantly different from the estimates for basic wages. In summary, this set of estimates shows that basic wages are substantially procyclical and extra pay components do not systematically respond to the business cycle. As such, these results suggest that the cyclicality of basic wages drives the cyclicality of firms' marginal labour costs.

There are some theoretical reasons to question whether the cyclicality of current ('spot') wages has any consequences for employment and output. When workers and firms form long-term employment relationships, current wages are better understood within a stream of payments in which year-to-year fluctuations might not be allocative for employment (Becker, 1962). Any stream of *remitted* wages that adds up to the same present value should not affect a firm's employment decisions. Instead, what matters to a firm when deciding to continue an employment relationship is whether the expected present value of a worker's output exceeds the expected present value of her labour costs. To understand the typical effects of changes in current wages on expected labour costs, we follow Grigsby, Hurst, and Yildirmaz (2021) and estimate the persistence of each pay component for job stayers. Intuitively, the more persistent a pay component is, the larger will be its impact on the present value of labour costs. In Appendix C, we show that receiving a certain extra pay component in one year significantly increases the likelihood of receiving this same pay component again the next year, within the same job. However, relatively high amounts of extra pay earned in one year tend to be followed by relatively low amounts in the next year. In contrast, if the current basic wage was high relative to the stream of past and future basic wages within a job, then a worker expects next year's basic wage also to be relatively high. This implies that a rise in basic wages generally leads to an increase in labour costs exceeding the initial value of that rise, while the opposite is true for extra pay. Combining the evidence presented in this and the preceding section, the representative payroll data from Great Britain demonstrate that basic wages are the most relevant measure of remuneration for macroeconomic models that rely on wage rigidities to generate muted responses of a firm's marginal labour costs to the business cycle.

4. The Extent of Downward Nominal Wage Rigidity

This section explores the data on job-stayer basic wages for evidence that nominal wages are rigid downwards. We begin by analysing commonly used statistical indicators of downward nominal wage rigidity. Subsequently, we apply insights from the theoretical framework developed by Elsby (2009) to investigate empirically whether downward rigidity constrained firms' nominal wage setting.

4.1. Statistical Indicators

To give a first impression of the possible extent of downward nominal wage rigidity, we construct histograms of year-to-year nominal log basic wage changes. Figure 2 displays these distributions for hourly-paid and salaried job stayers, pooled across all years in the sample period. Four key characteristics are visible. First, large numbers of job stayers experience year-to-year basic wage freezes: 21 percent of hourly-paid workers and 17 percent of salaried job workers receive no basic wage change. The height of these spikes at zero is striking and shows that wage freezes occur substantially more frequently than previously thought. Nickell and Quintini (2003) and Elsby, Shin, and Solon (2016) documented comparable spikes never exceeding 10 percent in Britain, though they were unable to study basic wages and did not account for some probable, but previously undocumented, sources of measurement errors in how payroll records were transferred into the British datasets. In Appendix D, we describe these sources of measurement errors, which would lead to an underestimation of year-to-year nominal wage freezes and an overestimation of wage cuts if not addressed. Second, the distribution of basic wage changes is markedly asymmetric, dropping off sharply below zero. Although the sample period spans a severe recession between 2007 and 2009, on average only 11 percent of salaried and 4 percent of hourly-paid job stayers experienced a year-to-year decrease in nominal basic wages. Third, the histogram shows a relative lack of small positive wage changes, directly to the right of zero. This would be consistent with the presence of adjustment costs, or 'menu costs', when firms adjust basic wages (e.g., Kahn, 1997). Fourth, the basic wages of hourly-paid job stayers show more signs of DNWR than for salaried employees; the spike at zero is 4 percentage points higher, and the share of cuts is 7 percentage points lower compared with the salaried job stayers.

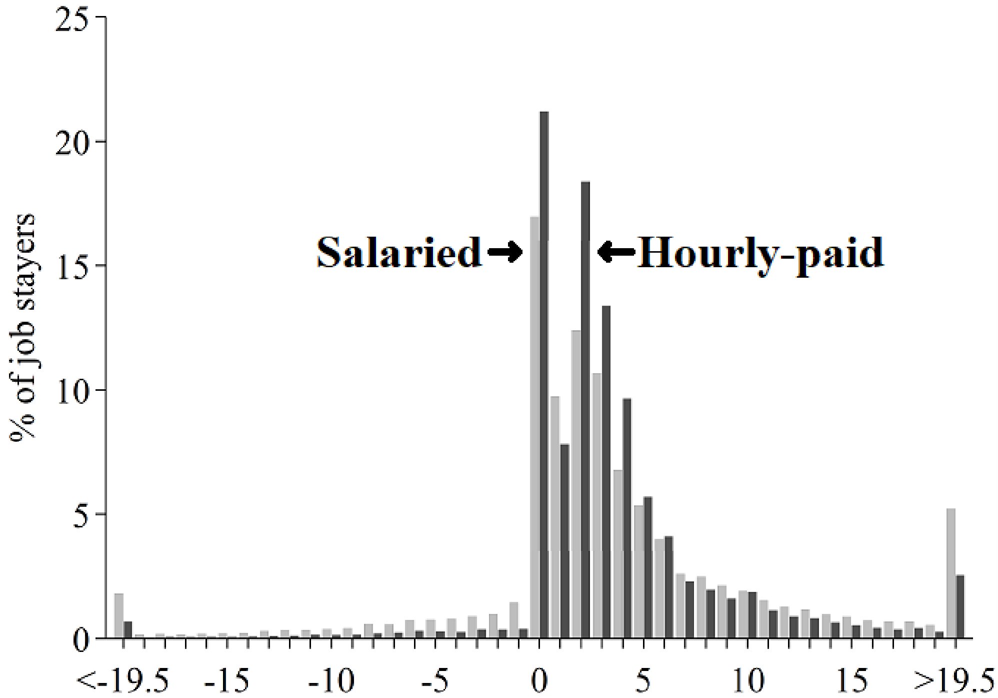

FIGURE 2: Frequency distribution of year-to-year changes in log basic wages for job stayers

Notes: The zero bin includes log changes in the interval $[-0.5, 0.5]$. The other bins are similarly one log point wide, with the positive changes excluding the lower limit and including the upper limit, and *vice versa* for negative changes. Data are pooled across all years in 2006-2018. **Light bars**: salaried. **Dark bars**: hourly-paid job.

Though striking, our discoveries that wage freezes are far more likely than previously thought and wage cuts are rare in Great Britain do not provide clear evidence for the presence of DNWR on their own. For example, if wages were set according to implicit contracts between firms and workers, then this could account for the observed spike without wages necessarily being subject to DNWR (Barro, 1977). Therefore, we consider a statistical indicator proposed by Lebow, Saks, and Wilson (2003) to assess whether DNWR was constraining wage cuts. This indicator, hereafter the 'LSW statistic', measures the asymmetry of the wage change distribution. Underlying the LSW statistic is the assumption that the notional wage change distribution is symmetric around its median, and that its upper half is largely unaffected by DNWR.[11] The LSW statistic is then calculated as the difference between the notional and observed mass of negative wage changes, and is interpreted as the missing mass of wage cuts caused by DNWR.[12] An increase in the LSW statistic indicates that this missing mass has increased, potentially providing evidence that DNWR has become more binding.

Figure 3 displays the LSW statistic for job stayers and each year in 2007-2018, as well as the UK unemployment rate and the inflation rate. For reference, we also show the share of year-to-year freezes among job-stayer wage changes over time. Wage freezes were

[11]The assumption that the notional wage change distribution is symmetric is commonly invoked, see, for example, Card and Hyslop (1996) or Dickens et al. (2007).

[12]The LSW statistic is computed as $[1-F(2\cdot P_{50}+0.005)]-F(-0.005)$, where $F(\cdot)$ is the cumulative distribution function of wage growth and P_{50} is its median. See Lebow, Saks, and Wilson (2003) for further discussion.

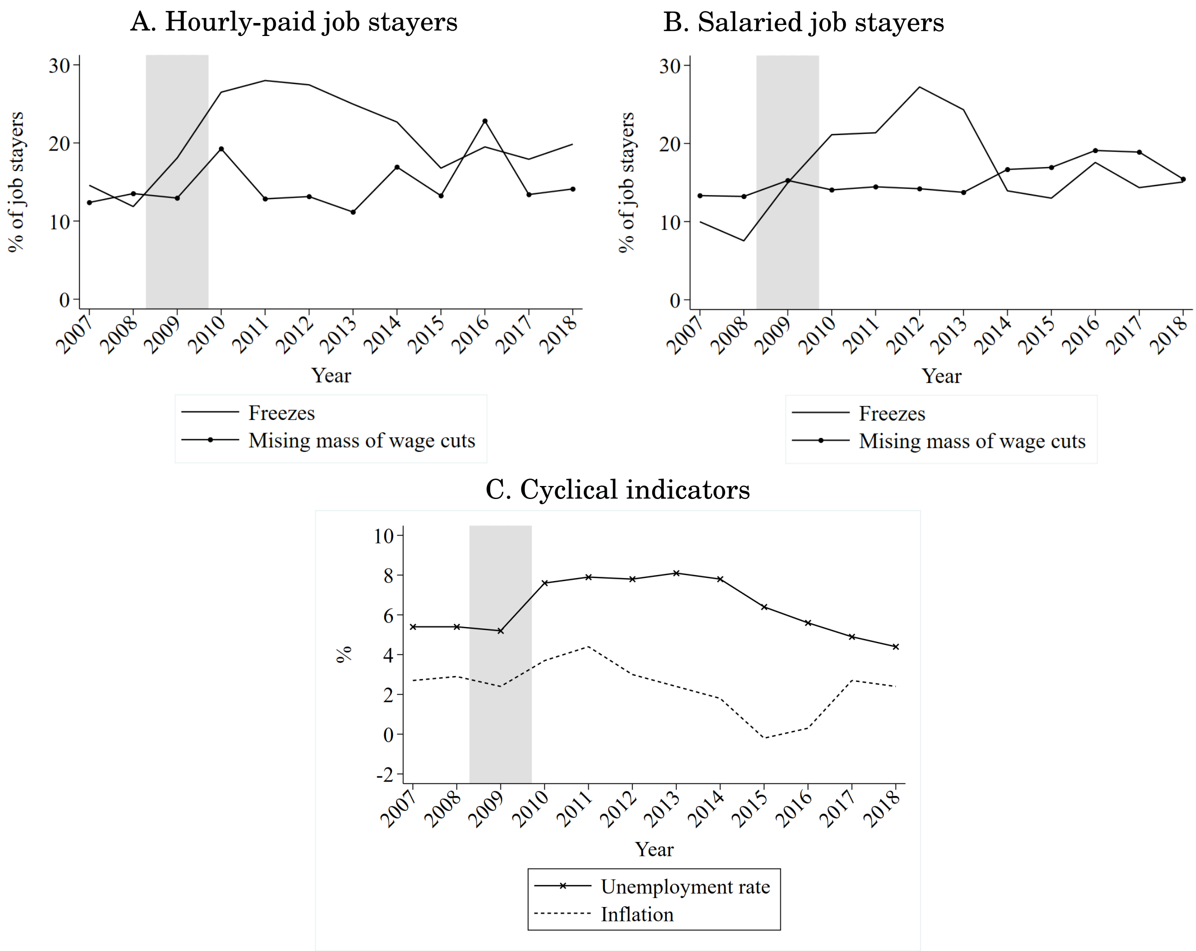

FIGURE 3: Statistical indicators of DNWR in the year-to-year basic wages of job stayers, and cyclical indicators

Notes: 'Freezes' show year-to-year changes in log basic wages for job stayers in the interval $[-0.5, 0.5]$. 'Missing mass of wage cuts' show the LSW statistic. Inflation is measured as the April-to-April log change in the UK Consumer Price Index (CPI). The unemployment rate refers to UK individuals aged 16 and over, seasonally adjusted and for the second April of each period, expressed as a percentage of the economically active population. Both series are from the Office for National Statistics. See Appendix Tables E1-E3 for the underlying statistics shown and sample sizes by year.

more frequent during the Great Recession and its aftermath. More than one-in-four job stayers experienced year-to-year basic wage freezes at the height of the recession, peaking at 28 percent and 27 percent of hourly-paid and salaried job stayers, respectively. The LSW statistic provides support for the presence of binding DNWR. The estimated missing mass of wage cuts increased during the recession to 15.3 percent for salaried job stayers and to 19.3 percent for hourly-paid job stayers, indicating that the wage change distribution became more asymmetric. The substantial spike in 2016 is due to the introduction of the National Living Wage to the UK.[13] The visible large spikes at zero and the lack of wage cuts in the histograms, combined with the cyclical increases in both year-to-year wage freezes and the

[13]In 2015, the UK government announced the introduction of a new National Living Wage that would apply to all workers aged 25 and over from April 2016. The rate was set at £7.20, which represented a substantial increase of 7.5 percent over the previous National Minimum Wage.

missing mass of wage cuts, provide statistical evidence that basic wages exhibit significant downward nominal rigidity in Great Britain.

Our result that the distribution of basic wage changes is asymmetric, falling off sharply to the left of zero, confirms the findings from previous studies on basic wage changes in other countries, regardless of the data source: Iceland (Sigurdsson and Sigurdardottir, 2016), France (Le Bihan, Montores, and Heckel, 2012), and the US (Card and Hyslop, 1996; Elsby, Shin, and Solon, 2016; Barattieri, Basu, and Gottschalk, 2014; Grigsby, Hurst, and Yildirmaz, 2021). Our evidence contrasts with the results of two recent US studies on basic wages. First, Barattieri, Basu, and Gottschalk (2014) analysed the Survey of Income and Program Participation (SIPP), reporting basic wage freezes of 40-50 percent in the US from year-to-year. However, wages can be reported with rounding errors in household surveys, tending to bias upward the observed prevalence of nominal wage freezes (Smith, 2000; Elsby and Solon, 2019). Second, Grigsby, Hurst, and Yildirmaz (2021) used data from a large US payroll service provider and found that around a third of job stayers experience year-to-year basic wage freezes. This value is significantly higher than what we find in Britain. However, as Grigsby, Hurst, and Yildirmaz (2021) explain, their sample under-represents large companies that normally use in-house payroll systems. In Section 6, we provide evidence that large firms are substantially more likely to adjust the basic wages of their employees. Additionally, Grigsby, Hurst, and Yildirmaz (2021) winsorised wage observations that fell below the US federal minimum wage. Since that minimum wage remained unchanged over their sample period, this potentially swept wage cuts into wage freezes, which would increase estimates of nominal rigidity. Combined, these facts may have led Grigsby, Hurst, and Yildirmaz toward finding more basic wage freezes than exist in nationally representative data.

Although the previous Section 3 argued that basic wages are the most relevant wage variable for macroeconomic models of wage stickiness, for comparison to the previous literature, specifically Nickell and Quintini (2003) and Elsby, Shin, and Solon (2016), we repeat the above analysis for exactly the same pay measure used in those studies: average earnings per hour, excluding overtime. We find evidence that this broader pay measure does exhibit signs of DNWR among job stayers, albeit to a lesser degree than basic wages (see Appendix F for these results and discussion). A possible reason why our conclusion differs from the previous UK literature, even when analysing the same pay measure, is that the rate of inflation was relatively high in earlier sample periods, possibly hiding the signs of DNWR. Additionally, we correct for previously undocumented measurement errors in the earlier or same payroll datasets (see Appendix D), which likely biased the estimates of wage flexibility upwards.

4.2. Wage Growth Compression

So far we have presented statistical evidence on the extent of DNWR, based on the assumption that the upper right-hand tail of the notional wage change distribution is unaffected by downward rigidity. However, Elsby (2009) argued that downward rigidity can also affect that upper tail. Therefore, we go beyond purely statistical indicators, using the main insights and predictions derived from the theoretical framework of Elsby (2009), to investigate whether firms are constrained by DNWR.[14]

In Elsby's discrete-time, intertemporal framework, downward wage rigidity arises because a nominal wage cut leads to a sharp reduction in worker productivity. This assumption is based on Bewley (1999), who was told by employers that the main reason they did not cut nominal wages was their belief that cuts harm employee morale, a key determinant of worker productivity.[15] In this way, if a firm raises its nominal wage to increase productivity, but has to cut the wage by an equal amount in the future, then the result will be an overall loss in productivity. This means that nominal wage increases become partially irreversible. In an uncertain world, forward-looking firms would thus refrain from raising wages today, or at least be conservative when doing so, because otherwise there would be an increased likelihood of having to cut wages in the future.

Theoretically, when an idiosyncratic productivity shock hits an employer-employee match, the response of the nominal wage depends on the extent of DNWR. Without DNWR, the nominal wage change would equal the change in nominal productivity. With binding DNWR, the optimal wage change policy would instead take the form of a trigger strategy. In this case, large positive productivity shocks lead to wage increases, while large negative shocks lead to wage cuts. For intermediate values of the shock, there is a range of inaction where nominal wages remain unchanged. When shocks are i.i.d. across employer-employee matches, this region of inaction shows up in the nominal wage change distribution as a spike at zero. Another consequence of DNWR implied by this theory is that if nominal wage changes occur, then they will be compressed relative to a world without DNWR.

Therefore, according to Elsby (2009), DNWR can be detected by observing the differential effects of inflation and productivity growth on the percentiles of the nominal wage growth distribution. In the absence of DNWR, *real* wages should move one-to-one with productivity growth and inflation should have no effect. In contrast, if DNWR constrained firms' wage setting, the theory predicts three effects: (1) A firm constrained by DNWR will moderate nominal wage raises, because raising wages today increases the likelihood of having to cut

[14] In a related study, Stüber and Beissinger (2012) also used the insights from Elsby (2009) to analyse real wage changes in West Germany.

[15] The explanation of Bewley (1999) for DNWR is not the only one possible. For example, MacLeod and Malcomson (1993) and Holden (1994) argue that past wages are the baseline while worker and firm are negotiating a new level of pay.

wages at a cost in the future. As inflation and productivity growth increase, the frictionless likelihood that the firm wishes to cut wages declines, and so the firm compresses nominal wage growth less. Therefore, on average, we should observe the upper percentiles of the real wage change distribution rising with inflation, and rising more than one-for-one with productivity. (2) Optimal wage setting under DNWR implies that there exists a range of inaction where nominal wages are kept constant. Equivalently, real wage growth equals minus the inflation rate. This is where inflation can *"grease the wheels of the labor market"* (Tobin, 1972), by bringing down the real labour cost without cutting nominal wages. (3) In very low percentiles of the wage growth distribution, nominal wage cuts will be compressed because of the disproportionate fall in productivity that they would cause. Higher productivity growth and/or inflation will increase wage growth in this lowest range, because firms expect that they will likely be able to reverse nominal cuts in the future and so refrain from making them in the first place. These effects should lead to a more than one-for-one increase in wage growth with productivity growth in the lowest percentiles. Table 3 summarises these predicted effects of inflation and productivity growth on the real wage growth distribution.

TABLE 3: Predicted effects of the inflation rate and productivity growth on the percentiles of the real basic wage growth distribution, according to Elsby (2009)

n-th percentile of the real wage growth distribution (P_n)	Coefficient on	
	Inflation rate	Productivity growth
$P_n >$ minus inflation rate	> 0	> 1
$P_n \approx$ minus inflation rate	< 0	< 1 (attenuates toward zero)
$P_n <$ minus inflation rate	> 0	> 1

To estimate the effects of inflation and productivity growth across the real wage growth distribution, we apply an unconditional quantile regression (UQR) approach (Firpo, Fortin, and Lemieux, 2009). As Stüber and Beissinger (2012) explain, this approach is preferable to the seemingly unrelated regressions (SUR) adopted by Elsby (2009) to test his theory, as UQR allows the whole distribution of the chosen set of explanatory variables to be taken into account when modelling real wage growth. To maintain consistency with the statistical evidence presented above, we estimate the UQR separately for salaried and hourly-paid job stayers. The dependent variables are the job-stayer re-centred influence functions of each respective year's log real basic wage changes. Our explanatory variables are the national inflation rate, measured as the change in the log of the UK Consumer Price Index (April-to-April), and a measure of regional productivity. Adapting the setup in Elsby (2009), we use average annual EU-NUTS1 regional log real wage growth among job stayers to proxy this productivity growth. Current and one-year lagged regional unemployment rates are included in the model as control variables, as well as region fixed effects. At the job-year level, we include further control variables for the gender, age, relative pay-level and collective agreement status of the employee, plus the private sector status and industry of the employer, as well as its size and growth in terms of the number of employees. The latter set of control variables is also used in the heterogeneity analysis later (Section 6) and are further described in Appendices A & G.

The results from estimating these UQRs are reported in Table 4 for selected percentiles, displaying only coefficients for the explanatory variables of interest, inflation and productivity growth, as well as the means of the dependent variable. These results provide clear evidence that the upper tail of the wage growth distribution in Great Britain is compressed by DNWR, for both hourly-paid and salaried job stayers. The estimated coefficients on inflation and productivity growth are consistent with the theoretical predictions in Elsby (2009) (Table 3). Specifically, at the 80-90th percentiles, the influence of inflation is significantly positive and the coefficients on productivity growth are substantially larger than one.

Table 4 also displays the estimates for the impact of productivity growth and inflation on the lower percentiles of the wage growth distribution. The theoretical predictions for these percentiles depend on the position of zero real wage growth in the overall distribution, which varies year-to-year and with the inflation rate. Over our whole sample period, annual average CPI inflation in the UK was 2.4 percent. As predicted by the theory, we find the most negative coefficient on inflation at the 20th percentile, and the coefficient on productivity growth is the most attenuated toward zero at this percentile. The results for the 10th percentile are also consistent with the theory, because relative to the 20th percentile the effect of inflation is diminished and the effect of productivity growth is larger. The evidence is somewhat mixed for intermediate percentiles of the wage growth distribution. Wage growth above the rate of inflation is predicted to increase with higher inflation rates, and the coefficient on productivity growth should theoretically be above one. However, annual inflation rates fluctuated from -0.2 to 4.4 percent over 2006-2018. These values imply that the 10-50th percentiles of the real wage growth distribution were at some point equal to minus the rate of inflation, for which the theory predicts a negative coefficient on inflation and an attenuation toward zero for the coefficient on productivity growth. Taken together, our findings provide significant evidence for compression across the basic wage growth distribution, and hence for the presence of binding DNWR.

5. Internal Wage Structures: Does Job-Stayer Wage Rigidity Matter?

Our analysis up to this point has focused on the basic wages of job stayers. But according to a large literature on labour markets with search frictions, the key determinant of a firm's hiring decision, and consequently aggregate unemployment fluctuations, is the flexibility of hiring wages (Pissarides, 2009). Even in this context, our findings regarding job-stayer wage changes remain relevant. Recent empirical studies confirm strong preferences for fairness of pay among employees; job satisfaction is found to depend on relative wage comparisons,

TABLE 4: The effects of inflation and productivity growth on percentiles of real basic wage growth for job stayers

	Salaried			Hourly-paid		
	Wage growth	Inflation	Productivity	Wage growth	Inflation	Productivity
	(I)	(II)	(III)	(IV)	(V)	(VI)
p10	−0.047	−0.265	0.364	−0.029	−0.261	0.321
		(0.026)	(0.028)		(0.011)	(0.010)
p20	−0.027	−0.391	0.240	−0.023	−0.618	0.086
		(0.008)	(0.007)		(0.011)	(0.012)
p30	−0.017	−0.300	0.771	−0.014	−0.409	1.022
		(0.011)	(0.013)		(0.021)	(0.024)
p40	−0.005	−0.164	0.878	−0.004	−0.206	0.738
		(0.009)	(0.010)		(0.013)	(0.012)
p50	0.003	−0.171	0.733	0.002	−0.350	0.522
		(0.011)	(0.010)		(0.015)	(0.015)
p60	0.013	−0.213	0.825	0.010	−0.508	0.646
		(0.011)	(0.009)		(0.026)	(0.024)
p70	0.027	0.043	0.938	0.019	−0.309	0.840
		(0.017)	(0.015)		(0.024)	(0.022)
p80	0.056	0.353	1.223	0.033	0.462	1.291
		(0.036)	(0.031)		(0.035)	(0.043)
p90	0.116	0.767	1.634	0.073	1.587	3.009
		(0.064)	(0.057)		(0.103)	(0.102)

Notes: Results of unconditional quantile regression, bootstrapped standard errors in parentheses, 50 replications. Gaussian kernel and Silverman's plug-in optimal bandwidth.
Sample sizes: Salaried job stayers, 697,480; hourly-paid job stayers, 303,336.
Controls included: regional unemployment rate and its lag, indicator variables for NUTS1 British regions, as well as the control variables included in the probit model estimates described in Section 6 & Appendix G, except the indicator for a rounded hourly rate of pay.
The explanatory variables are 'Inflation': measured as the change in the log of the UK CPI (April-to-April); and 'Productivity': the average real wage growth each year in the corresponding region for salaried or hourly-paid job stayers.

and wage inequality lowers retention rates and productivity (Card et al., 2012; Breza, Kaur, and Shamdasani, 2018; Dube, Giuliano, and Leonard, 2019). These ideas are captured in the theoretical framework of Snell and Thomas (2010). They show that it can be optimal for a firm to commit to offering the same wage contracts to new hires and existing workers. This occurs when risk neutral firms insure their risk averse employees against income fluctuations, and this insurance motive creates rigidity for existing employees that can prevent the hiring wage from falling sufficiently during recessions, causing cyclical unemployment fluctuations.

Empirically, if firms were constrained by internal wage structures as described above, then hiring wages should be as cyclical as the wages of existing employees. To assess the evidence for internal wage structures in Great Britain, we adapt the empirical strategy of Snell, Stüber, and Thomas (2018). The estimation is carried out in two steps. In the first step, basic wages are regressed on year-region dummy variables and a set of worker-firm-level controls, distinguishing between new hires and existing employees. New hires are all workers who joined the firm within the given year, and all other workers are incumbents. As explained by Snell, Stüber, and Thomas, it is important to control for as much worker-firm heterogeneity as possible when analysing internal wage structures. Failing to control for cyclical changes in match quality - the average quality of a match formed during an expansion is usually considered to be higher (Hagedorn and Manovskii, 2013) - may lead to an overestimation of the cyclicality of hiring wages. Therefore, we include match fixed effects (MFEs) and proxies for worker-firm tenure in the first-step regression to account for cyclical match quality. One potential downside of this approach is that any difference in the levels of the hiring and incumbent wages, that is constant over the duration of a match, will be absorbed by the MFEs.[16] In the second step, we regress the sets of estimated year-region wage effects from the first step on business cyclical indicators, thus generating separate estimates for the cyclicality of new-hire and incumbent-worker wages. The key test is whether these cyclicality estimates differ significantly between new hires and incumbents.

The first-step regression is given by:

$$\log(\tilde{w}_{ijrt}) = \theta_{ij} + \mathbf{1}\{\text{Incumbent}_{it}\} \cdot \beta^{I}_{rt} + \mathbf{1}\{\text{New hire}_{it}\} \cdot \beta^{N}_{rt} + \tilde{\mathbf{x}}'_{it}\tilde{\boldsymbol{\delta}} + \tilde{\varepsilon}_{ijrt} \ , \tag{2}$$

where $\tilde{w}_{ijrt}$ is the basic wage of worker i, in firm j, which is in region r, and in year t. The firm-worker-match fixed effect is θ_{ij}, and $\tilde{\mathbf{x}}_{it}$ is a vector of time-varying worker characteristics. This vector includes a worker's age and age squared in years, and worker-firm tenure squared. The indicator variable $\mathbf{1}\{\text{Incumbent}_{it}\}$ equals one if worker i is an incumbent in year t and zero otherwise. Similarly, $\mathbf{1}\{\text{New hire}_{it}\}$ equals one if the wage observation belongs to a new hire. The coefficient estimates of interest are $\hat{\beta}^{I}_{rt}$ and $\hat{\beta}^{N}_{rt}$, which give the composition-adjusted period means, within each region, of the incumbent and hiring wages, respectively.

In the second step, we regress changes in the estimated β-coefficients on cyclical indicators:

$$\Delta\hat{\beta}^{z}_{rt} = c^{z}_{u} + \gamma^{z}_{u}\ U_{rt} + \eta^{z}_{rt} \ , \tag{3}$$

$$\Delta\hat{\beta}^{z}_{rt} = c^{z}_{y} + \gamma^{z}_{y}\ \Delta y_{rt} + e^{z}_{rt} \ , \tag{4}$$

for $z = \{I, N\}$. The business cycle indicators considered in region r and period t are the unemployment rate U_{rt}, and the change in the logarithm of nominal gross value added (GVA),

[16] The first-step controls also for cyclical changes in the average worker quality, which is another important source of cyclical composition bias (Solon, Barsky, and Parker, 1994).

Δy_{rt}. The estimates of γ_u^z measure the semi-elasticity of nominal wage growth with respect to the unemployment rate of incumbents ($z = I$) and new hires ($z = N$). Similarly, the estimates of γ_y^z show the elasticity of incumbent wage growth with respect to nominal GVA growth. Wages may be correlated over time within a region, and also across regions in a given period. Therefore, we cluster standard errors in the second step at both the region and year levels.

Table 5 displays the estimates of γ_u^z and γ_y^z from the second step, i.e., Equations (3) & (4), where the first-step Equation (2) was estimated separately for hourly-paid and salaried workers. The year-to-year growth in the basic wages of incumbent workers decreases significantly by 0.36 percent when the unemployment rate increases by one percentage point. The wages of new hires generally show stronger responses: basic wage growth declines by 0.44 percent for salaried workers and by 0.52 percent for hourly-paid workers when the unemployment rate increases by one percentage point. These differences appear to be not economically important. However, to assess whether hiring wages are statistically more cyclical than incumbent wages, we compute the differences between the estimated year-region coefficients from the first step, $(\hat{\beta}_{rt}^I - \hat{\beta}_{rt}^N)$, and regress these on the cyclical indicators in the second step. The results are shown in column (III) of Table 5. The coefficients are precisely estimated but not significantly different from zero, suggesting that new-hire and incumbent wage growth exhibit the same cyclicality with respect to the unemployment rate. Importantly, a similar set of conclusions applies when we use regional gross value added instead of the unemployment rate as the business cycle indicator in Equation (4) (rows 3-4, Table 5). We also carry out these estimations for earnings per hour to see whether firms use extra pay components to circumvent internal wage structures. The results are very similar to those for basic wages (Appendix Table F1). Specifically, there is no economically or statistically significant difference in the cyclicality of incumbent and new-hire wage growth over the business cycle. These findings support the notion that firms are constrained when setting hiring wages by their internal wage structures.

A challenge to the interpretation of the estimation results in Table 5 is that employees are typically hired into long-term employment relationships. As argued by Kudlyak (2014), it is therefore not the *spot* wage of new hires that matters for hiring decisions, but rather the 'user cost of labour'. These user costs are defined as the present expected value from hiring a worker in the current period compared to waiting until the next period. However, as Gertler and Trigari (2020) explained, this user cost argument implicitly assumes that the composition of new hires does not depend on the state of the business cycle. When new hires and existing workers receive equal wages, conditional on match-fixed effects and worker-skills, then the current wage represents the user cost. In this way, incumbent wages offer a better guide to the cyclicality of the user cost of labour, instead of using hiring wages which the researcher might only be able to adjust incompletely for the cyclical changes in the composition of new hires. Grigsby, Hurst, and Yildirmaz (2021) also found that hiring and

TABLE 5: Estimated basic wage responses of new hires and job stayers to regional unemployment and gross value added

Dependent variable / Cyclical indicator	Incumbents (≥ 1 year in job) (I)	New hires (< 1 year in job) (II)	Difference (Stayer-hire) (III)
Unemployment rate:			
1. Basic wages, salaried	−0.0036	−0.0044	0.0008
	(0.0014)	(0.0015)	(0.0015)
2. Basic wages, hourly-paid	−0.0036	−0.0052	0.0016
	(0.0011)	(0.0012)	(0.0012)
Nominal GVA:			
3. Basic wages, salaried	1.033	0.853	0.180
	(0.192)	(0.740)	(0.678)
4. Basic wages, hourly-paid	0.720	0.375	0.346
	(0.156)	(0.413)	(0.363)

Notes: Standard errors in parentheses robust to two-way clustering over year and EU-NUTS1 region.
Sample sizes: first-step estimates for salaried basic wages 1,079,062, for hourly pay rate 754,578. Second-step estimates based on 121 year-region observations.
Column (I) shows estimates of the cyclical response of wages from the second-step regressions for job stayer wages. Column (II) shows the equivalent estimates for new-hire wages. Column (III) shows estimates from second-step regressions where the dependent variable is the difference between the incumbent and new hire coefficients from the first step, i.e., $\hat{\beta}^I_{rt} - \hat{\beta}^N_{rt}$.
Rows 1-2 show semi-elasticity estimates of basic wages with respect to the unemployment rate.
Rows 3-4 show elasticity estimates of basic wages with respect to gross value added.
Sources: Wages are from the Annual survey of Hours and Earnings. NUTS1 unemployment rates are from the Office for National Statistics (ONS) for April of each year, corresponding to the reporting period of the ASHE. Regional GVA are also from the ONS and correspond to calendar years, i.e., April-to-April changes in wages are regressed on the annual change in GVA over the years prior to each April.

incumbent wages are equally cyclical in the US. Hence, studying job-stayer wages is sufficient to uncover meaningful estimates for the business cycle responses of the user cost of labour.

6. Heterogeneity in Wage Freezes and Cuts across Workers and Firms

The ASHE data contain detailed job characteristics reported by employers (e.g., collective pay agreements) or retrieved from administrative data sources (e.g., number of employees in the firm). To isolate the effects of various worker, job, and firm characteristics on the probability that a year-to-year basic wage cut or freeze is observed, we estimate probit models separately for salaried and hourly-paid job stayers. Appendix G contains full details and results, and here we provide a brief summary, focusing on the most salient patterns.

We find statistically significant gender differences in conditional year-to-year basic wage changes. Cuts are more common for job stayers who are male rather than female, and basic wages freezes are more commonly found among hourly-paid men (Tables G1 and G2, first and second rows). However, the magnitudes of these differences do not seem economically significant for the average job stayers.[17] Private sector job stayers are less likely to experience wage freezes than non-private sector workers. Job stayers in a firm where wages are set according to a collective agreement are on average significantly less likely to see their wages cut or frozen. We find stark age-specific differences in the conditional probabilities of having basic wages cut or frozen: Job stayers aged 15-29 years are significantly less likely to receive a wage freeze than 45-64 year-olds (salaried, 9.6 percentage points less; hourly-paid, 7.2 percentage points less), and their probabilities of basic wage cuts are also lower. These results are surprising, because we control for firm growth and the place of workers in the earnings distribution, which should mostly control for worker life-cycle effects. Job stayers with earnings below 2/3 of the median are both significantly less likely to receive the same wage next year and significantly more likely to receive a basic wage cut, compared to median or high earners. The model estimates also show that the probability of a wage cut is significantly higher outside the "Wholesale & Retail Trade, Hotels & Restaurants" sectors, and basic wage freezes for hourly-paid job stayers occur significantly more frequently (5.5 percentage points) outside this sector.

Focusing on one of the starker sets of differences between jobs, Figure 4 displays the distribution of job-stayer log changes in basic wages separately for small firms with less than 50 employees, medium-sized firms who employ between 50 and 249 workers, and large firms with 250 employees or more. As rows 13-16 in Table G1 show, wage freezes in very large British firms with more than 5,000 employees are strikingly less common than in smaller firms, conditional on several other observable characteristics of job stayers. The conditional probability of basic wage freezes for salaried employees in small firms is 11.8 percentage points greater than in very large firms, and 19.1 percentage points greater for hourly-paid job stayers. This is mainly accounted for by less common moderate basic wage growth in small firms, rather than fewer wage cuts in small firms, as Figure 4 shows. These results support the US-based findings of Kurmann and McEntarfer (2019), who documented similar features for the distribution of changes in total earnings per hour (basic wages plus extra pay) by firm size. The US payroll dataset studied by Grigsby, Hurst, and Yildirmaz (2021) substantially under-represents very large firms with more than 5,000 employees, which can account for a sizable portion of employment - the median firm size in the British payroll data for job stayers is over 2,500 employees (see Table 1). Grigsby, Hurst, and Yildirmaz acknowledged this as a limitation in their study. Taken together, these results suggest that the higher year-to-year zero basic wage change spikes reported by Grigsby, Hurst, and Yildirmaz, of around a third of firm stayers, could be accounted partially by their lack of coverage of very large firms.

After controlling for other observable differences between job stayers, wage freezes are significantly more likely in shrinking than expanding firms, approximately by 5 percentage points (rows 19-21, Tables G1 & G2). These results support the same inverse relationship

[17] At the sample average and conditional on other characteristics, salaried (hourly-paid) male job stayers are 0.008 (0.014) more likely to have their wages cut than salaried (hourly-paid) female job stayers.

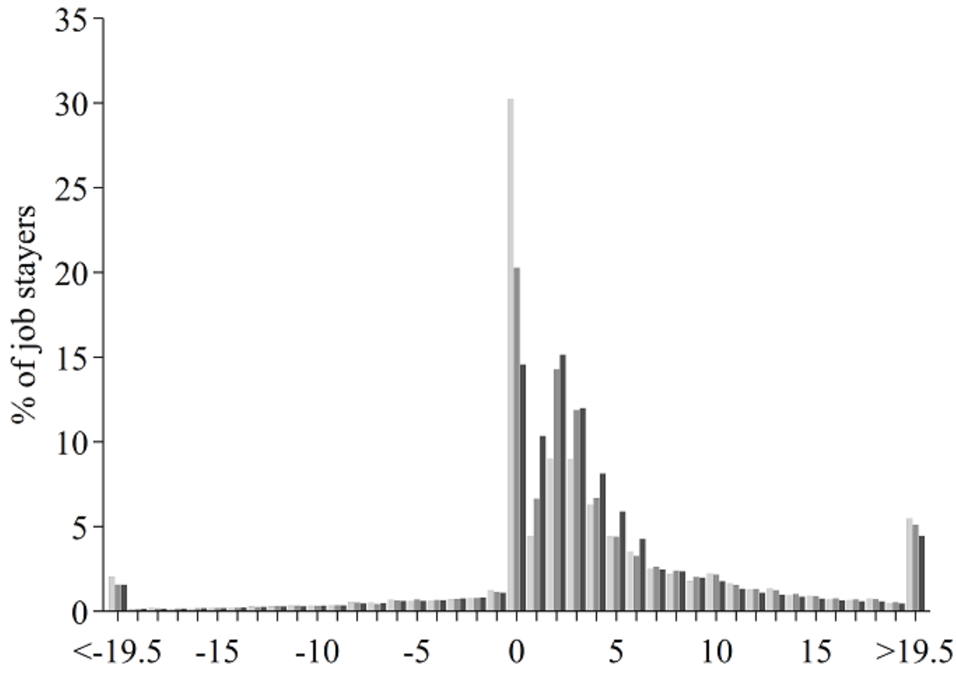

FIGURE 4: Frequency distribution of year-to-year changes in log basic wages of salaried and hourly-paid job stayers by firm size.

Notes: See Figure 2. Pooled data for 2006-18. **Light grey bars**: small firms (< 50 employees). **Medium grey bars**: medium-sized firms (50 – 249 employees). **Dark grey bars**: large firms (≥ 250 employees). The size of a job stayer's firm is defined in the first of two consecutive years.

between the frequency of total earnings per hour freezes and firm growth found by Kurmann and McEntarfer (2019) in the US. However, our data do not allow us to assess empirically how many workers in shrinking firms lost their jobs because of the nominal rigidity in basic wages, since we cannot exclude that negative idiosyncratic shocks to firm-level productivity caused both more layoffs and more freezes when comparing shrinking and expanding firms.[18]

Lastly, we focus on hourly-paid workers, for whom employers directly report basic wage rates in the ASHE. We find that year-to-year freezes are significantly more common among hourly-paid job stayers who had basic wages in the previous year that were multiples of ten pence (rows 22-23, Table G1).[19] There are two possible explanations. First, employers are incorrectly reporting rounded basic wages. Second, round basic wages are accurately reported and just happen to be more rigid, for example, if round basic wages are somehow preferred. This idea can be found in the price-point literature, where a set of pre-specified

[18]We also find that the results are not notably different when we focus on firms with more substantial employment changes of at least 10% or 1000 employees. Further, see Appendix Figure E1 for the distribution of year-to-year log basic wage changes, dependent on whether a firm's total number of employees was shrinking or expanding between the same years. A relatively large zero spike for shrinking firms is clearly visible, and positive wage growth is more likely in expanding firms. Note, job stayers in 2007-08 and 2017-18 are excluded from this analysis because it appears as though the ONS imputed the number of employees in firms using the previous years' values.

[19]Kahn (1997), using data on household heads from the Panel Study of Income Dynamics in 1970-1988, documented that hourly pay rates ending in exact dollar or half-dollar amounts made up 50 percent of the job stayers with hourly rate freezes.

prices, or wages in our case, can simplify the decision problem of boundedly rational workers and employers. For example, Hahn and Marenčák (2020) have shown that incorporating output-price points can improve how New Keynesian models match key business cycle statistics, such as the dynamics of the inflation rate. The extent to which 'wage-points' can improve the performance of monetary models might be a promising area for future research.

7. Macroeconomic Consequences of Downward Nominal Wage Rigidity

What are the macroeconomic implications of our main findings? In the New Keynesian model, the workhorse model of monetary policy analysis, wage rigidity is the crucial nominal friction required to match the persistence of inflation and output, and for monetary policy to have real effects (Christiano, Eichenbaum, and Evans, 2005). This class of models normally assumes that nominal wages can only be adjusted every period with a constant probability, so-called 'Calvo wage setting' (Calvo, 1983). Our findings that basic wage freezes affect on average 17 percent of salaried job stayers and 21 percent of hourly-paid job stayers implies constant quarterly wage-change probabilities of 0.36 and 0.32, respectively.[20] These values support the degree of nominal wage rigidity that is required as an assumption in New Keynesian models to fit the observed persistence of output, inflation, and unemployment in the US (Christiano, Eichenbaum, and Evans, 2005) and the UK (Faccini, Millard, and Zanetti, 2013).

7.1. Simulating the Welfare Costs of Downward Nominal Wage Rigidity and Low Inflation

Our estimates for the extent of DNWR in Great Britain have potentially sizeable output consequences. We demonstrate this using the theoretical framework of the Phillips curve from the dynamic stochastic general equilibrium model of Benigno and Ricci (2011) (hereafter BR).[21] In their model, the long-run inflation-output trade off caused by DNWR is the result of nominal rigidities that distort wage adjustments. Aggregate and idiosyncratic shocks lead to intratemporal and intertemporal shifts of real wages and employment across sectors. If a negative demand shock hits a firm, then it would like to reduce its real wages. When inflation is low, firms are more likely to need to decrease the nominal wage to bring about the desired real wage cut. With binding DNWR, nominal wages cannot be cut and, as a consequence, firms will instead reduce their employment and output. By allowing real wages to fall sufficiently, moderate inflation in the model can grease the wheels of the labour market.

[20] The values are converted as follows: $p_q = 1-(1-p_y)^{1/4}$, with p_q denoting the quarterly wage-change probability, and p_y the yearly wage-change probability.

[21] We use and adapt the model's replication code available on the website of the *American Economic Review*. Daly and Hobijn (2014) develop a similar framework of the Phillips curve, assuming a different wage setting process proposed by Benigno and Ricci (2011), however, this does not affect the model's long-run outcomes.

In BR's model, households enjoy the consumption of differentiated goods, which are supplied by firms operating in competitive output markets. The prices of consumption goods are fully flexible. Labour markets are characterised by monopolistic competition, in which households supply a variety of labour to the output producing firms. The output of each firm is produced using a constant elasticity of substitution aggregate production function, with labour as the only input. This production function exhibits the 'love-for-variety' property, such that every household supplies its labour to all firms. Forward-looking firms and workers internalise the effect that their current wage setting will have on the likelihood of being constrained by downward rigidity in the future, consistent with Elsby (2009); firms which raise wages today are more likely to be bound by downward nominal rigidity in the future. Therefore, firms which experience a positive demand shock will raise wages less than they would have done in the absence of nominal rigidities. DNWR thus leads to wage growth moderation. However, despite this forward-looking wage setting, firms are sometimes hit by negative demand shocks that are sufficiently large, given the current level of their real wages, and firms are forced to decrease employment.

First, we calibrate the version of BR's model presented in section V.C of their article, wherein nominal wages cannot be cut, except when a 'very large' shock hits the firm. This captures the finding by Bewley (1999) that workers are more likely to agree to wage cuts when the survival of the entire firm is at stake. Specifically, we choose the frequency at which such very large idiosyncratic productivity shocks hit firms to match the shares of basic wage freezes and cuts that we observe in British payroll data (Section 4), keeping all other parameters at the values calibrated by BR based on the US. We set the model parameter governing the frequency of large shocks to $\lambda dt = 0.05$, implying that a firm's wages are expected to become downward flexible one quarter in every five years. Figure 5 displays the frequency distribution of year-to-year log wage changes simulated by the model. It generates frequencies of wage freezes (18.7 percent) and wage cuts (9.1 percent), which correspond closely to our main empirical findings. Additionally, the model replicates the empirical feature of a markedly asymmetric wage-change distribution, with positive changes occurring far more frequently than negative changes.

Figure 6 shows the model-generated long-run relationship between the average rate of inflation (vertical axis) and the equilibrium output gap (in percent).[22] We plot this simulated long-run Phillips curve for two different frequencies at which wages are assumed to become flexible. The line representing $\lambda dt = 0$ shows the extreme scenario in which nominal wages can never be cut. The second line represents our calibrated value of $\lambda dt = 0.05$. First, note that DNWR implies an important change from the standard result with flexible wages: the long-run Phillips curve is no longer vertical. For both values of λdt, the long-run output loss is around 0.2 percent of GDP, relative to an economy with flexible wages, when the annual rate of wage inflation is 4 percent. When the average rate of wage inflation decreases to 2 (1)

[22] We impose the constraint that the long-run equilibrium employment level cannot exceed the population size.

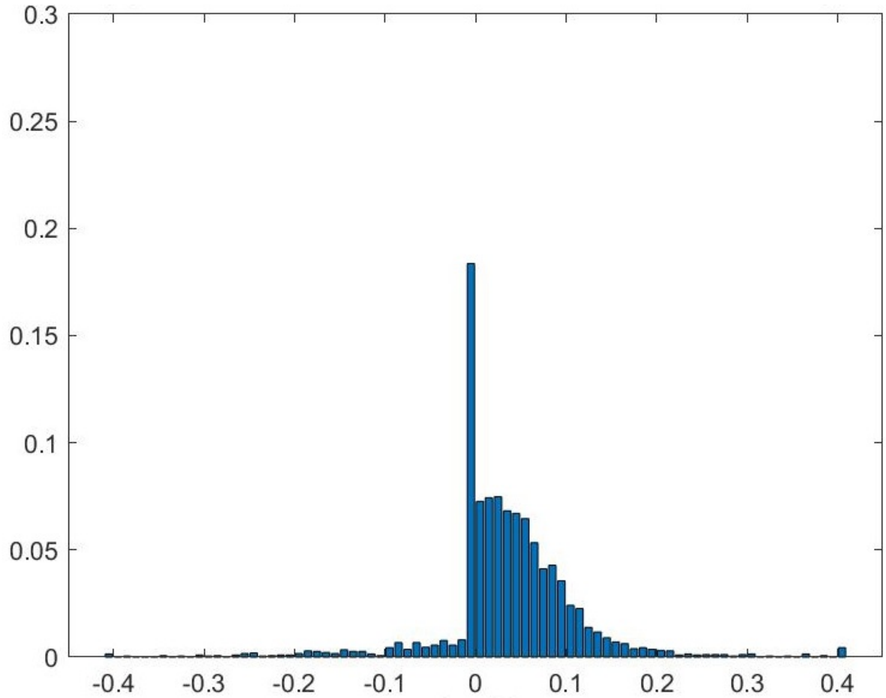

FIGURE 5: Simulation results for year-to-year changes in log wages

Notes: Equivalent to Benigno and Ricci (2011): Figure 5, p. 1458. Simulation uses exactly the same parameters as BR, except $\lambda dt = 0.05$, meaning nominal wages become flexible one quarter in every five years, on average.

percent, the output gap becomes more substantial: for our calibrated parameter of $\lambda dt = 0.05$, the equilibrium output loss amounts to around 0.7 (1.3) percent of GDP. When the mean inflation rate is low, the constraint on downward nominal wage changes is more likely to bind and be more costly in terms of higher unemployment.

As a robustness check, we conduct another simulation exercise. We consider an economy where only some workers are subject to DNWR, while the wages of the other workers can be adjusted flexibly. We use the model presented in section V.B of BR, no longer allowing for the previous large idiosyncratic shocks. We parameterise the share of workers with rigid wages, α, to the observed share of workers in our data for which basic wages make up their entire labour income: 53 percent. Figure 7 displays the long-run Phillips curve of this economy. Even if DNWR only affects around half of all workers, the resulting negative output gap is sizeable, with the model implying that around 0.9 percent of long-run equilibrium output is lost when mean annual inflation is 2 percent.

In summary, the simulations of BR's model, which match our empirical findings, suggest that DNWR causes a long-run negative output gap of around 0.7 to 1.3 percent of GDP in the prevailing low-inflation environment in Great Britain. These substantial economic costs point to an optimal rate of inflation that may be significantly positive rather than zero or negative (the Friedman rule). For example, increasing the equilibrium rate of inflation from 1 to 4 percent per year would reduce the long-run output loss from 1.3 to 0.2 percent of GDP.

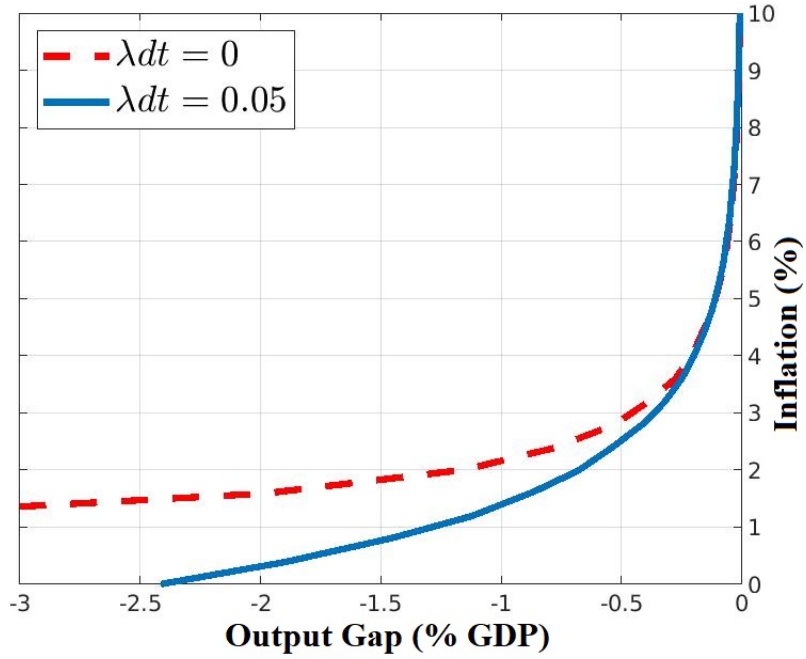

FIGURE 6: Long-run Phillips curve with binding DNWR and very large infrequent idiosyncratic shocks

Notes: Equivalent to Benigno and Ricci (2011): Figure 6, p. 1459. Simulation uses the same parameter values as BR, except λ. The dashed line shows the long-run output gap (in percent of GDP) when nominal wages can never be cut. The solid line shows the gap when nominal wages become downward flexible once every five years, on average.

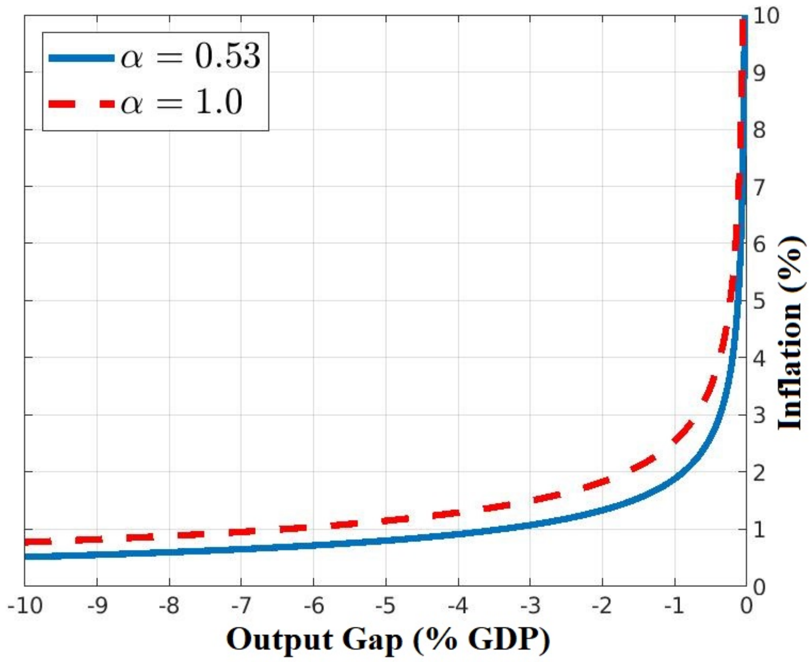

FIGURE 7: Long-run Phillips curve with binding DNWR for an α-share of workers

Notes: Equivalent to Benigno and Ricci (2011): Figure 4, p. 1455. Simulation uses the same parameter values as BR, except α. The dashed line shows the long-run output gap (in percent of GDP) when nominal wages can never be cut. The solid line shows the gap when nominal wages can never be cut for an α-share or employees, while the remaining $(1-\alpha)$-share of employees has completely flexible nominal wages.

8. Conclusion

Stubbornly low inflation in many developed countries has pushed the topic of DNWR back to the centre stage of macroeconomics. In this paper, we use unique longitudinal employer-employee data from Great Britain, the Annual Survey of Hours and Earnings (ASHE), to analyse the extent of DNWR. These payroll-based data allow us to accurately distinguish between basic wages and extra pay components, such as overtime and incentive pay. Since a wealth of recent findings from administrative data has led researchers to question whether DNWR is actually a pervasive and consequential feature of labour markets, the need for additional evidence on this subject is great.

We provide several new findings. First, basic wages, which we argue are the relevant wage measure from a macroeconomic perspective, are downward rigid. Among employees who are constantly employed in the same job from year-to-year, every fifth job stayer, on average, receives a constant basic wage and this figure increased to around 28 percent during the Great Recession. Basic wage cuts occur far less frequently than previously thought, with 11 percent of salaried job stayers and only 4 percent of hourly-paid job stayers receiving them on average each year. We show that the distribution of wage changes is markedly asymmetric, and provide strong evidence for the compression of wage growth in continuing employment relationships, suggesting that firms are constrained by binding DNWR.

Our findings support the anecdotal evidence presented by Bewley (1999), who gathered information on nominal pay setting by interviewing US employers. Bewley found that employers were reluctant to cut nominal wages because they feared it would damage worker morale, and in this way it would negatively impact productivity, labour turnover, and the recruitment of new employees. However, Bewley also found that most managers believed that cutting incumbent workers' wages would not prevent layoffs; labour is only a small share of variable costs and the short-run price-elasticity of product demand is low, such that pay cuts would create little extra work.[23] Although it is not clear how this anecdotal evidence generalises to the aggregate economy, it highlights the need to explore channels other than layoffs through which DNWR can have real effects.

Connecting to a large literature on labour markets with search frictions, our second contribution shows that the wages of new hires are just as responsive to the business cycle as the wages of incumbent workers. This suggests that internal wage structures are important (Snell and Thomas, 2010), and provides a channel through which DNWR in job-stayer wages can cause cyclical unemployment fluctuations. Our third contribution demonstrates substantial heterogeneity in the apparent extent of DNWR, conditional on observable firm and worker-characteristics. Most notably, we find a strong negative correlation between firm

[23] Employers that responded to financial distress during a recession with pay cuts reported to Bewley only minor problems with worker morale and productivity. However, it is possible that those employers that faced the least adverse consequences were exactly the ones enacting pay cuts.

size and the share of employees that experiences basic wage freezes. Job stayers in firms with less than 50 employees are around twice as likely to receive the same wage year-to-year as job stayers in very large firms with more than 5,000 employees. This might partially explain the main differences between our results and the recent US study by Grigsby, Hurst, and Yildirmaz (2021), who may have documented more rigidity in basic wages because their dataset under-represented very large firms.

Our findings have important implications for macroeconomic models and monetary policy. We find empirical support in representative payroll microdata for the degree of nominal rigidity typically required in New Keynesian models to match the persistence of output and inflation (Christiano, Eichenbaum, and Evans, 2005). Additionally, our results justify the assumption of downward nominal rigidity invoked in recent macroeconomic models of business cycle fluctuations (e.g., Benigno and Ricci, 2011; Daly and Hobijn, 2014; Dupraz, Nakamura, and Steinsson, 2019). Researchers who are investigating macroeconomic variables in a low-inflation environment should consider incorporating DNWR into their models. Our result that basic wage freezes are far more common when wage rates are in multiples of ten pence suggests a potentially fruitful area for future research. Akin to the literature on price-points, where a set of pre-specified prices can simplify the decision problem of boundedly rational agents and improve the match of New Keynesian models to key business cycle statistics (Hahn and Marenčák, 2020), wage-points could improve the performance of monetary models further. Related, there are other patterns in these data which merit further thought and study, such as the fact that basic wage freezes are far more common in smaller firms.

9. References

Akerlof, George A., William T. Dickens, and George L. Perry. 1996. "The Macroeconomics of Low Inflation." *Brookings Papers on Economic Activity*, 27(1): 1–59.

Altonji, Joseph G., and Paul J. Devereux. 2000. "The extent and consequences of downward nominal wage rigidity." In *Research in Labor Economics*. 383–431. Emerald Group Publishing Limited.

Barattieri, Alessandro, Susanto Basu, and Peter Gottschalk. 2014. "Some Evidence on the Importance of Sticky Wages." *American Economic Journal: Macroeconomics*, 6(1): 70–101.

Barro, Robert J. 1977. "Long-term contracting, sticky prices, and monetary policy." *Journal of Monetary Economics*, 3(3): 305–316.

Becker, Gary S. 1962. "Investment in Human Capital: A Theoretical Analysis." *Journal of Political Economy*, 70(5): 9–49.

Bell, David N. F., and Robert A. Hart. 2003. "Wages, Hours, and Overtime Premia: Evidence from the British Labor Market." *Industrial and Labor Relations Review*, 56(3): 470–480.

Benigno, Pierpaolo, and Luca A. Ricci. 2011. "The Inflation-Output Trade-Off with Downward Wage Rigidities." *American Economic Review*, 101(4): 1436–1466.

Bewley, Truman F. 1999. *Why wages don't fall during a recession.* Harvard University Press.

Breza, Emily, Supreet Kaur, and Yogita Shamdasani. 2018. "The Morale Effects of Pay Inequality." *Quarterly Journal of Economics*, 133(2): 611–663.

Calvo, Guillermo. 1983. "Staggered prices in a utility-maximizing framework." *Journal of Monetary Economics*, 12(3): 383–398.

Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez. 2012. "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction." *American Economic Review*, 102(6): 2981–3003.

Card, David, and Dean Hyslop. 1996. "Does Inflation "Grease the Wheels of the Labor Market"?" NBER Working Paper Series 5538.

Carneiro, Anabela, Pedro Portugal, and Jose Varejao. 2014. "Catastrophic job Destruction during the Portuguese Economic Crisis." *Journal of Macroeconomics*, 39(B): 444–457.

Christiano, L., M. Eichenbaum, and G. Evans. 2005. "Nominal rigidities and the dynamic effects of a shock to monetary policy." *Journal of Political Economy*, 113(1): 1–45.

Daly, Mary C., and Bart Hobijn. 2014. "Downward Nominal Wage Rigidities Bend the Phillips Curve." *Journal of Money, Credit and Banking*, 46(2): 51–93.

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

Dickens, William T., Lorenz Goette, Erica L. Groshen, Steinar Holden, Julian Messina, Mark E. Schweitzer, Jarkko Turunen, and Melanie E. Ward. 2007. "How Wages Change: Micro Evidence from the International Wage Flexibility Project." *Journal of Economic Perspectives*, 21(2): 195–214.

Dube, Arindrajit, Laura Giuliano, and Jonathan Leonard. 2019. "Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior." *American Economic Review*, 109(2): 620–663.

Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson. 2019. "A Plucking Model of Business Cycles." NBER Working Paper 26351.

Elsby, Michael W. L. 2009. "Evaluating the economic significance of downward nominal wage rigidity." *Journal of Monetary Economics*, 56(2): 154–169.

Elsby, Michael W. L., and Gary Solon. 2019. "How Prevalent Is Downward Rigidity in Nominal Wages? International Evidence from Payroll Records and Pay Slips." *Journal of Economic Perspectives*, 33(3): 185–201.

Elsby, Michael W. L., Donggyun Shin, and Gary 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.

Faccini, Renato, Stephen Millard, and Francesco Zanetti. 2013. "Wage Rigidities in an Estimated Dynamic, Stochastic, General Equilibrium Model of the UK Labour Market." *Manchester School*, 81(S1): 66–99.

Fallick, Bruce, Daniel Villar, and William Wascher. 2020. "Downward Nominal Wage Rigidity in the United States during and after the Great Recession." Federal Reserve Bank of Cleveland Working Paper 16-02R.

Fehr, Ernst, and Lorenz Goette. 2005. "Robustness and real consequences of nominal wage rigidity." *Journal of Monetary Economics*, 52(4): 779–804.

Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux. 2009. "Unconditional Quantile Regression." *Econometrica*, 77(3): 953–973.

Gertler, Mark Huckfeldt, Christopher, and Antonella Trigari. 2020. "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires." *Review of Economic Studies*, 87(4): 1876–1914.

Grigsby, John, Erik Hurst, and Ahu Yildirmaz. 2021. "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data." *American Economic Review*, 111(2): 428–471.

Hagedorn, Marcus, and Iourii Manovskii. 2013. "Job Selection and Wages over the Business Cycle." *American Economic Review*, 103(2): 771–803.

Hahn, Volker, and Michal Marenčák. 2020. "Price points and price dynamics." *Journal of Monetary Economics*, 115(C): 127–144.

Hazell, Jonathon, and Bledi Taska. 2020. "Downward Rigidity in the Wage for New Hires." CFM Discussion Paper Series 2020-28.

Holden, Steinar. 1994. "Wage bargaining and nominal rigidities." *European Economic Review*, 38(5): 1021–1039.

Jardim, Ekaterina S., Gary Solon, and Jacob L. Vigdor. 2019. "How prevalent is downward rigidity in nominal wages? Evidence from payroll records in Washington State." NBER Working Paper Series 25470.

Kahn, Shulamit. 1997. "Evidence of Nominal Wage Stickiness from Microdata." *American Economic Review*, 87(5): 993–1008.

Keynes, John. M. 1936. *The General Theory of Employment, Interest and Money.* Palgrave Macmillan.

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

Kurmann, André, and Erika McEntarfer. 2019. "Downward Nominal Wage Rigidity in the United States: New Evidence from Worker-Firm Linked Data." Center for Economic Studies, U.S. Census Bureau Working Papers 19-07.

Le Bihan, Herve, Jeremi Montores, and Thomas Heckel. 2012. "Sticky Wages: Evidence from Quarterly Microeconomic data." *American Economic Journal: Macroeconomics*, 4(3): 1–32.

Lebow, David. E., Raven. E. Saks, and Beth. A. Wilson. 2003. "Downward Nominal Wage Rigidity: Evidence from the Employment Cost Index." *Advances in Macroeconomics*, 3(1): 1–28.

MacLeod, W. B., and J. M. Malcomson. 1993. "Investments, Holdup, and the Form of Market Contract." *American Economic Review*, 83(4): 811–837.

McLaughlin, Kenneth. J. 1994. "Rigid wages?" *Journal of Monetary Economics*, 34(3): 383–414.

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

Office for National Statistics. 2005. "The new Annual Survey of Hours and Earnings questionnaire."

Office for National Statistics. 2019. "Annual Survey of Hours and Earnings, 1997-2018: Secure Access. [data collection]." 14th Edition. UK Data Service. SN: 6689, http://doi.org/10.5255/UKDA-SN-6689-13.

Pissarides, Christopher A. 2009. "The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?" *Econometrica*, 77(5): 1339–1369.

Schaefer, Daniel, and Carl Singleton. 2019. "Cyclical Labor Costs Within Jobs." *European Economic Review*, 120.

Schaefer, Daniel, and Carl Singleton. 2020*a*. "Nominal Wage Adjustments and the Composition of Pay: New Evidence from Payroll Data." Department of Economics, Johannes Kepler University Linz, Austria Economics working papers 2020-11.

Schaefer, Daniel, and Carl Singleton. 2020*b*. "Recent changes in British wage inequality: Evidence from large firms and occupations." *Scottish Journal of Political Economy*, 67(1): 100–125.

Schmitt-Grohé, Stephanie, and Martin Uribe. 2016. "Downward Nominal Wage Rigidity, Currency Pegs, and Involuntary Unemployment." *Journal of Political Economy*, 124(5): 1466–1514.

Sigurdsson, Josef, and Rannveig Sigurdardottir. 2016. "Time-dependent or state-dependent wage-setting? Evidence from periods of macroeconomic instability." *Journal of Monetary Economics*, 78(C): 50–66.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review*, 97(3): 586–606.

Smith, Jennifer C. 2000. "Nominal Wage Rigidity in the United Kingdom." *Economic Journal*, 110(462): 176–195.

Snell, Andy, and Jonathan P. Thomas. 2010. "Labor Contracts, Equal Treatment, and Wage-Unemployment Dynamics." *American Economic Journal: Macroeconomics*, 2(3): 98–127.

Snell, Andy, Heiko Stüber, and Jonathan P. Thomas. 2018. "Downward real wage rigidity and equal treatment wage contracts: Theory and evidence." *Review of Economic Dynamics*, 30(3): 265–284.

Solon, Gary, Robert Barsky, and Jonathan A. Parker. 1994. "Measuring the Cyclicality of Real Wages: How Important is Composition Bias?" *Quarterly Journal of Economics*, 109(1): 1–25.

Stüber, Heiko, and Thomas Beissinger. 2012. "Does downward nominal wage rigidity dampen wage increases?" *European Economic Review*, 56(4): 870–887.

Tobin, James. 1972. "Inflation and Unemployment." *American Economic Review*, 61(1): 1–18.

The Extent of Downward Nominal Wage Rigidity: New Evidence from Payroll Data

Online Appendix

Daniel Schaefer Carl Singleton[†]

Appendix A. Further Description of the Data

In what follows, we provide further details on the datasets used. All the relevant documentation and variable descriptions are publicly available from the UK Data Service. The ONS has also published various documents concerning the quality and consistency of the ASHE.

For each year in the period 2004-2018, the ASHE should in principle be a random sample of all UK employees, irrespective of occupation, size of employer, etc. Given the legal obligation of employers to respond using their payrolls, it has a high response rate and is generally considered to be very accurate. Conditional on a hundred percent response and success in locating employees and employers, the ASHE should be a true one percent random sample of UK employees in each year: all those with an administrative life-long National Insurance number which has a numerical part ending in two specific digits are included in the sampling frame every year, making it a worker-level panel dataset. However, there are two major sources of systematic under-sampling of employees, both occurring if individuals do not have a current tax record when they are looked up in January each year. This could happen for some individuals who have recently moved jobs, or for those who earn too little (mostly part-time) to pay income tax or National Insurance. From 2004, the ASHE, after it replaced its predecessor the New Earnings Survey, aimed to sample some of those employees which were likely under-represented before. It added supplementary responses for those without income-tax payments, and also attempted to track employees whose jobs changed between the determination of the sampling frame in January and the reference period in April. 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. One exception to this occurs where individuals have been assigned temporary National Insurance numbers, typically in the case of non-UK nationals. Following ONS recommendations, we drop all person observations where there are inconsistencies in age or sex within the National Insurance number-based ASHE longitudinal person identifier variable, *piden*.

The ASHE data contain information on the legal status of a firm, obtained from the administrative Inter-Departmental Business Register (IDBR). We classify private companies, sole proprietors, and partnerships as belonging to the private sector, while state-owned enterprises, nationalised industries, central government, local authorities, and non-profit organisations form the non-private sector in our analysis. The measure of firm size is also taken from the IDBR and refers to the total number of employees in the enterprise. An employee is working full-time hours when she works at least thirty hours per week. Some businesses have an arrangement with the ONS to provide their data electronically, which seems to be the case for very large enterprises in particular. Employees of such enterprises tend to have lower basic wages. Whenever it is reported that an employee's basic wage is

[†]email: daniel.schaefer@jku.at; corresponding author. This work is based on the Annual Survey of Hours and Earnings Dataset (Crown copyright 2019), 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.

set with reference to an agreement affecting more than one worker, e.g., pay agreed by a trade union or a workers' committee, we classify this employee's wage as affected by a collective agreement.

From 2005, a new questionnaire was introduced to the ASHE, which was intended to reduce the latitude for respondents' own interpretations of what was being asked of them. Prior to 2005, job stayers were identified by asking whether the job title and the description were the same as in the last reference period. From 2005 onward, the question became more specific, asking whether the employee had worked in the same job and role for more than a year. This relates to the ASHE identifier variable for job stayers, *sjd*, which we use to define job stayers for our analysis. For a very small number of job stayers, apparent inconsistencies can arise in the dataset: a worker can be marked as being in the same job, but the enterprise reference number can change. This is probably because of a change in a firm's ownership structure (e.g., a merger or a succession), which would result in a new administrative enterprise reference number being assigned, while the ASHE questionnaire explicitly tells the respondent to ignore such changes when answering whether an employee is working in the same job. We also ensure that each job stayer in our analysis features consecutive year-to-year job and wage records.

To avoid any potential errors, we also drop any worker observations for years with multiple job holdings, because in these cases it is not clear to which job the same-job marker refers to in the following year. Although the ASHE dataset is available from 2004 onward, we use only data starting from 2006 because some questionnaire changes in 2005 and 2006 introduced inconsistencies between these earlier years. Specifically, new instructions were included on how firms should report employees' hours worked. Before applying any further sample selection, the 2006-2018 ASHE panel dataset contains 2,096,927 worker-year observations. From this, we keep only observations where an employee is aged 16-64, and which have not been marked as having incurred a loss of pay in a 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, dropping trainees and apprenticeships. This is practically the same filter applied by the ONS in their published results on UK "Patterns of Pay" using the ASHE. We drop observations with missing basic hours or earnings variables. We drop observations with over a hundred or less than one basic hour worked, as these could reflect measurement error or the inclusion of overtime. Applying these steps leaves 1,924,873 worker-year observations. To further address some potential for measurement error, we drop a further 3,571 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 behaviour which has been hypothesized by the ONS. Note, any such rounding for pay measures, conditional on accurate hours records, would tend to bias upwards the incidence of nominal wage rigidity. The ASHE has introduced some earnings imputations, using similar matched 'donor' observations where responses were, for example, missing an entry of basic hours but had recorded basic pay. Hourly pay rates were never imputed. We exclude from our main results the approximate one percent of the remaining observations where imputations were made to hours or basic pay, in the period they can be identified of 2013-2018. By focusing on this period, we were able to confirm that not dropping these observations did not meaningfully alter any of our results. Finally, before focusing on job stayers, we trim the dataset by sequentially dropping the top and bottom 1% of basic wage observations. We trim twice as there are a tiny number of observations where there are clear coding errors for basic wage changes, such as £9.34 per hour in one year but £934 in the next year for a job stayer.

We use the following pay variables from the ASHE, summarised in Table A1. Basic pay is a worker's regular pay before adding anything extra. The dataset contains two measures of labour input: basic hours and overtime, whereby the former also includes hours worked at a shift premium. If an employer calculates an employee's basic pay by multiplying their basic hours worked by a pay rate, then the employer is also required to report this rate. The employer also reports the pay period, which is most often either weekly, fortnightly, four-weekly or a calendar month. We do not observe the reported totals for these periods, but instead the ONS derives weekly averages of variables.

TABLE A1: Overview of pay variables

		Description
	ASHE variables	
1.	Basic pay	All basic pay, excluding any extra payment
2.	Basic hours	Hours relating to basic pay (incl. hours paid at shift premium)
3.	Hourly pay rate	Reported hourly rate of pay where applicable (see text)
4.	Overtime pay	Total overtime pay in reference period
5.	Shift premium pay	Premium payments for shift work, night or weekend work
6.	Incentive pay	Incentive pay received for work carried out in the pay period
7.	Other pay	Pay received for other reasons, e.g., travel allowances
8.	Gross pay	Total pay received
9.	Annual gross pay	Annual gross earnings paid for the preceding tax year
10.	Annual incentive pay	Component of annual gross pay from incentive payments
	Derived variables	
11.	Basic wage	Basic pay divided by basic hours
12.	Earnings per hour, excl. overtime	Gross pay, excluding overtime pay, divided by basic hours
13.	Gross earnings per hour	Gross pay divided by the sum of basic and overtime hours

The key earnings variables that we analyse are the answers to the following questions in the ASHE questionnaire, whereby monetary values are measured in Pound sterling (GBP), including pence, and time variables are reported in hours and minutes:

BPAY: "How much basic pay, before deductions, did the employee receive in the pay period?
Include: all basic pay, relating to the pay period, before deductions for PAYE, National Insurance, pension schemes, student loan repayments and voluntary deductions. Include paid leave (holiday pay), maternity/paternity pay, sick pay and area allowances (e.g., London).
Exclude: pay for a different pay period, shift premium pay, bonus or incentive pay, overtime pay, expenses and the value of salary sacrifice schemes and benefits in kind."

BHR: "How many basic hours does [basic pay] relate to?
If your pay period is calendar month and hours are weekly, multiply the weekly hours by 4.348 to get calendar month hours. If the employee uses a decimal clock, please convert to hours and minutes. For example, 4.3 hours should be 4 hours and (0.3 multiplied by 60) minutes = 4 hours 18 minutes.

Include: any hours paid at shift premium and paid hours even if not worked.
Exclude: any hours paid as overtime."

HPAY: "[W]as the employee's basic pay in the pay period calculated by multiplying the number of hours they worked by an hourly rate of pay? [If yes,] what was the employee's hourly rate of pay in the pay period?"

OVPAY: "How much overtime pay did the employee receive for work carried out in the pay period?
Exclude: any basic, shift premium and bonus or incentive pay in this period, as well as overtime pay from the previous pay period."

SPPAY: "How much shift premium pay did the employee receive in the pay period?
Include: the element of shift premium pay. For example, for a 35 hour pay period, if the basic rate is £10 per hour and the premium rate is £12 per hour, multiply the difference of £2 by the hours worked (i.e. 35 multiplied by 2). The shift premium pay reported would therefore be £70.
Exclude: any basic, overtime and bonus or incentive pay."

IPAYIN: "How much [bonus or incentive payments did the employee receive,] related to work carried out in the pay period?
For example, if [an annual bonus was paid], the value should be divided by 12 if the employee was paid on a calendar month basis.
Include: profit sharing, productivity, performance and other bonus or incentive pay, piecework and commission.
Exclude: basic, overtime and shift premium pay."

OTHPAY: "How much pay did the employee receive for other reasons in the pay period?
Include: for example, car allowances paid through the payroll, on call and standby allowances, clothing, first aider or fire fighter allowances.
Exclude: paid leave (holiday pay), basic, overtime, shift premium, maternity/paternity, sick, bonus or incentive pay, redundancy, arrears of pay, tax credits, profit share and expenses."

GPAY: "How much gross pay, before deductions, did the employee receive for work carried out in the pay period?
Include: pay before deductions for PAYE, National Insurance, pension schemes, student loan repayments and voluntary deductions. Include basic, overtime, shift premium, bonus or incentive pay and any other pay.
Exclude: expenses and the value of salary sacrifice schemes."

AGP: "How much annual gross pay did the employee receive in their current job?
Include: pay before deductions for PAYE, National Insurance, pension schemes and voluntary deductions. Include basic, overtime, shift premium, profit sharing, productivity performance and bonus or incentive pay.
Exclude: any payments for expenses or previous employment."

ANIPAY: "How much of [AGP] is related to bonus or incentive payments for their current job?
Include: profit sharing, productivity performance and other bonus or incentive pay, piecework

and commission.

Exclude: basic, overtime and shift premium pay."

One take-away from these questions is that basic pay excludes shift premium pay, while basic hours include shift premium hours. Therefore, if the shift pattern of a worker changes between years, basic hours might change while basic pay remains unchanged, all else equal.

For reference, Table A2 display the distribution of observation over industries, separately for salaried and hourly-paid job stayers.

TABLE A2: Distribution of job stayers over industry sectors

Industry (SIC2003)	Hourly-paid (I)	Salaried (II)
A-F: Agric., Mining, Manuf., Energy, Constr. etc.	0.20	0.15
G-H: Wholesale & Retail Trade, Hotels & Restaurants	0.32	0.13
M-N: Education & Health	0.27	0.30
I-L, O: Other Services	0.22	0.42

Notes: Statistics use the later period for each job-stayer observation. Classification according to the ONS Standard Industrial Classification 2003. 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 Details on the Composition of Pay

This section provides more details on the extra pay components besides basic wages, combined for hourly-paid and salaried job stayers. Figure B1A shows that incentive pay, i.e., bonuses and commission for work carried out during the reference period in April, contributes less to total weekly earnings when moving up the basic wage distribution. When Nickell and Quintini (2003) studied incentive pay in Great Britain, they found that almost 22 percent of job stayers in the New Earnings Survey received such payments, compared with only 7 percent in our sample. This difference might be explained by a change in the incentive pay definition between their dataset and the ASHE. The new definition in the ASHE questionnaire, by focusing on incentive payments earned and paid in the April pay period, is more precise and gives more consistent estimates between years (Office for National Statistics, 2005). Hence, Figure B1A is likely to understate the relative importance of annual bonuses, because these are typically paid between January and March in the UK (Schaefer and Singleton, 2020*b*).

To assess the importance of bonus payments outside the April reference period, we compute the share of *annual* incentive pay in *annual* earnings. Both values refer to the preceding tax year, and so complete information about those variables should have been available to employers when the questionnaire was completed. Figure B1B shows that the share of incentive pay in annual earnings increases along the basic wage distribution. This contrasts with the results for incentive pay earned and received in April. Within the highest decile, the importance of annual incentive payments is greater, especially among the top percentiles. The differences between Figure B1A and Figure B1B most likely originate in the composition of incentive pay. While the left panel probably reflects a larger share of commission-type payments earned throughout the year, annual incentive pay captures bonus payments, particularly for salaried high-earners.

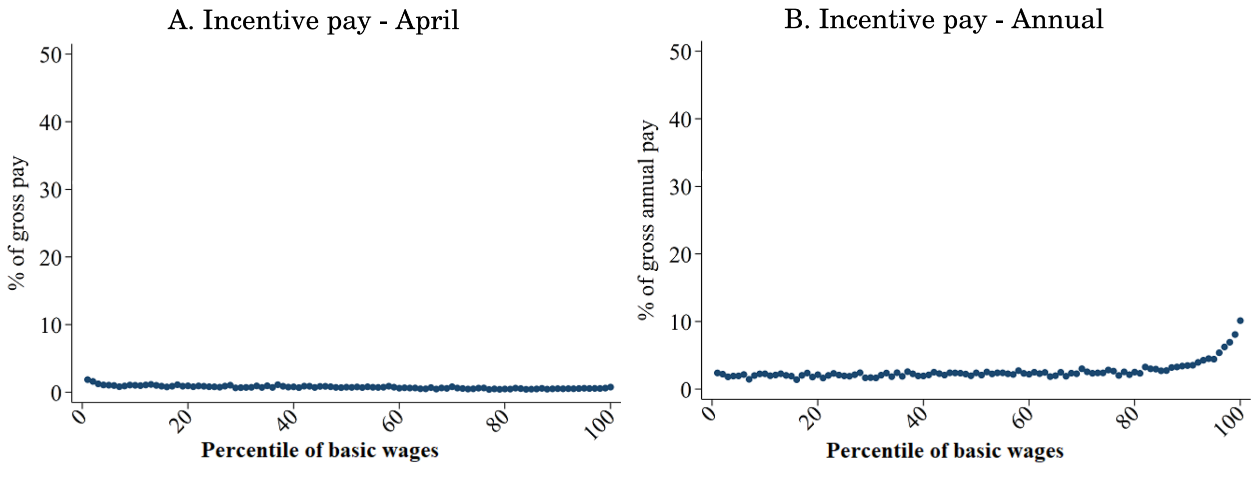

FIGURE B1: Incentive pay along the basic wage distribution

Notes: Panel A: Incentive pay in total earnings within the corresponding percentile of the basic wage distribution in April. Panel B: Average shares of annual incentive pay in annual earnings within the corresponding percentile of the basic wage distribution. Data pooled across all years.

The importance of paid overtime for job stayers declines with the basic wage (Figure B2B). While overtime accounts for almost 5 percent of total earnings in the bottom percentile and around 3 percent

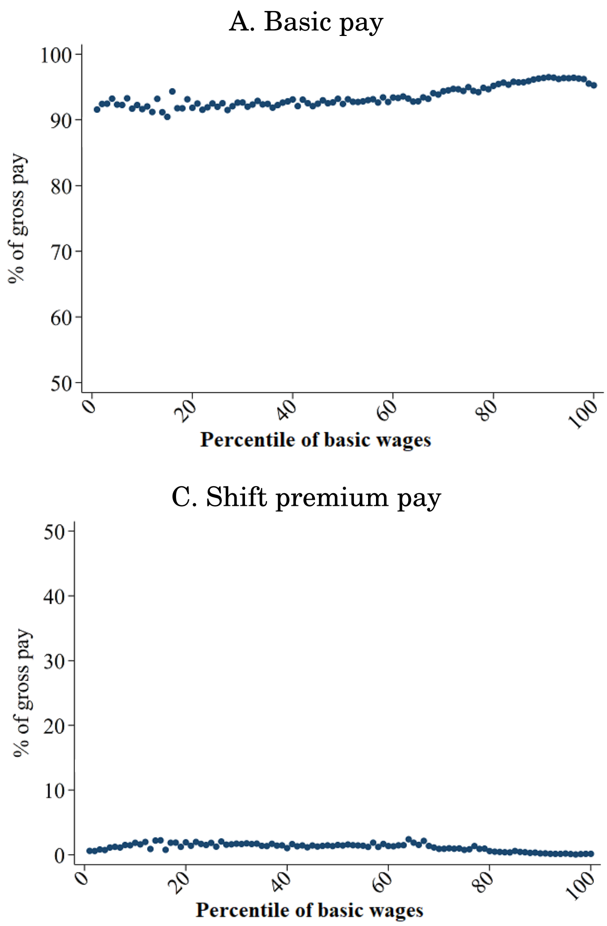

FIGURE B2: Non-incentive pay along the basic wage distribution

Notes: Average shares of pay component within the corresponding percentile of the basic wage distribution. Data pooled across all years.

at the median, its share in total earnings is less than 1 percent in the highest decile.[3] Shift premiums typically contribute less than 2 percent of total earnings, and are negligible in the top percentiles of basic wages (Figure B2C). A high basic wage is associated with a high share of total earnings from other pay, such as travel allowances (Figure B2D). Generally, overtime, shift premiums, and incentive pay all contribute relatively little to the level of total earnings in the top percentiles.

[3]See Bell and Hart (2003) for more details on the relationship between overtime hours and overtime premiums in Great Britain.

Appendix C. The Persistence of Pay Components

To assess the persistence of the different components of pay, we ask two questions about the extensive and intensive margin of extra pay. First, how likely is it that a job stayer receives some form of extra pay, conditional on her having received the same component last year? Second, conditional on receiving extra pay in consecutive years, how is the pay amount changing? To answer the first question, we estimate a linear probability model for the propensity to receive extra pay components:

$$D_{ijt} = \gamma D_{ijt-1} + u_{ijt} \tag{5}$$

where D_{ijt} equals one if worker i in job j in year t receives a particular extra pay component. We restrict the estimation sample to year-to-year job stayers. The results are reported in Table C1. Columns (I)-(IV) show that the propensity to receive non-basic pay in year t, conditional on having received the same extra pay in year $t-1$, is positive and significant. For example, salaried job stayers who received a positive amount of overtime pay last year, have a 52.3 percent likelihood of being paid for overtime in the current year, on average.

TABLE C1: Estimated year-to-year persistence of receiving extra pay components for job stayers

	Overtime (I)	Incentive pay (II)	Shift premium pay (III)	Other pay (IV)
Salaried	0.523^{***}	0.481^{***}	0.523^{***}	0.614^{***}
	(0.002)	(0.003)	(0.002)	(0.002)
Hourly-paid	0.546^{***}	0.539^{***}	0.546^{***}	0.522^{***}
	(0.002)	(0.004)	(0.002)	(0.002)

Notes: Least squares coefficient estimates of γ in Equation (5).
*** indicates significance from zero at the 0.1% level, two-sided tests, and standard errors in parentheses are robust to clustering within match.
Sample sizes: 691,521 salaried job-stayers and 416,497 hourly-paid job-stayers

For a better understanding of the persistence of individual pay components and, therefore, their likely impacts on the present value of labour costs, we follow Grigsby, Hurst, and Yildirmaz (2021) and estimate first-order autocorrelation coefficients at the intensive margin of each component, conditional on the receipt of the extra pay component in consecutive periods. For our sample of job stayers, we estimate the following regression, using least squares:

$$\log(\overline{w}_{ijt}) = \theta_{ij} + \rho \log(\overline{w}_{ijt-1}) + \overline{\mathbf{x}}_{it}'\overline{\boldsymbol{\delta}} + v_{ijt} \tag{6}$$

where $\overline{w}_{ijt}$ is the wage of worker i, in job j, and in year t; θ_{ij} is a fixed effect for an employer-employee match; ρ measures the first-order autocorrelation; and v_{ijt} is a random error term. The vector $\overline{\mathbf{x}}_{it}$ includes time-varying controls for firm size, age, age squared, and tenure squared.

Table C2 shows the autocorrelation estimates, computed separately for salaried and hourly-paid job stayers. The first column shows that basic wages are significantly positively autocorrelated. If basic wages are high in the current year, relative to the within-match average, then firms and workers can expect that basic wages will also be relatively high in the next year. This implies a strong impact of basic wage changes on the present value of the labour cost: for example, not only are the current basic wages high, they are also expected to exceed their average value in the match in subsequent years, conditional on the age of the employee and length of the match. The extra pay components affect the present value of labour costs less. For example, conditional on receiving overtime pay in two consecutive years, column (II) of Table C2 shows that the amount of overtime pay is significantly negatively autocorrelated across years within a match. This means that higher overtime pay than normal in one year is followed by lower overtime pay than normal in the next year, on average. This would serve as a strong offsetting force on the effect of overtime on the present value of labour costs. Similar results hold for incentive pay and shift premium pay (columns (III)-(IV)). However, for salaried job stayers, the autocorrelation of shift premium pay and other pay is not significantly different from zero at the 5 percent level, suggesting that, within a match, the amounts of these components are i.i.d., conditional on receipt. For hourly-paid job stayers, shift premium pay and other pay are significantly negatively autocorrelated. Taken together, these results suggest that basic wages are considerably more important for the present value of labour costs than extra pay components; the estimates here point to the allocative wage for job stayers in the ASHE being most closely approximated by basic wages.

TABLE C2: Estimated persistence of basic wages and non-basic pay components, job stayers

	Basic wages (I)	Overtime (II)	Incentive pay (III)	Shift premium pay (IV)	Other pay (V)
Salaried	0.220^{***}	-0.116^{***}	-0.068^{***}	-0.021	0.020
	(0.005)	(0.007)	(0.016)	(0.017)	(0.011)
Hourly-paid	0.199^{***}	-0.125^{***}	-0.129^{***}	-0.063^{***}	-0.064^{***}
	(0.008)	(0.007)	(0.021)	(0.012)	(0.012)
N (in 000s)					
Salaried	592	47	16	31	78
Hourly-Paid	325	54	11	35	29

Notes: Least squares coefficient estimates of ρ in Equation (6).
***, *, * indicate significance from zero at the 0.1%, 1%, 5% levels, respectively. Two-sided tests, and standard errors in parentheses are robust to clustering within match.
Sample for each pay component restricted to job stayers who received the pay component in consecutive years.

Appendix D. Measurement Error in Job-Stayer Wage Changes

This appendix discusses a previously undocumented source of errors in the ASHE dataset, concerning the accuracy of the reported hours of work. If an employee's pay period is a calendar month, but working hours are weekly, then the ASHE asks employers to multiply the weekly hours by 4.348 and to report the result as hours per calendar month. In our sample, almost 75 percent of job stayers are paid per calendar month. Because the questionnaire only allows employers to report hours and minutes worked, respondents have to round decimal values that result from the conversion of weekly hours. Unfortunately, no guidelines are provided to employers on how they should round in such cases.

For example, if an employee works 40 basic hours per week, then the hours worked in the calendar month are 173 hours, 55 minutes, 12 seconds ($40 \times 4.348 = 173.92$). It is in the discretion of the person who is answering the questionnaire whether she rounds up this number to 173 hours and 56 minutes, or rounds it down to 173 hours and 55 minutes. Another potential source of error is the conversion from decimal to minutes: the employer might incorrectly supply decimal numbers instead of hours and minutes. For example, a 38-hour week results in a monthly value of 165.224 ($38 \times 4.348 = 165.224$), so 165 hours and 13 minutes (omitting seconds). It is conceivable that some employers might incorrectly report 165 hours and 22 minutes.[4] The ONS converts these reported values back to average weekly hours, dividing monthly hours by 4.348. This can potentially explain the relatively high frequency of weekly hours worked in the ASHE dataset with values of 39.999 and 40.003.

To gain a better understanding of the significance of this novel source of measurement error, we compute the share of job stayers with non-zero changes in working hours of less than one minute from year-to-year. This is the case for around 21.5 percent of salaried job stayers who are paid per calendar month in consecutive years. To check whether these small hours changes are more likely to be erroneous or actual changes, we turn to these affected job stayers' *weekly* pay. These data should be reported directly from payroll, without requiring a conversion for the pay period, and so we expect them to be relatively more accurate. Only 6 percent of the affected job stayers show adjustments in weekly basic pay. This strongly suggests that the recorded marginal changes in weekly hours are more likely to be measurement error than actual changes.

We perform another check by looking at changes in weekly basic pay, which are unaffected by any measurement error in basic hours worked in the data. We consider *exact* zero changes, instead of a range around zero, to define a year-to-year freeze, because payroll data on weekly basic pay are more reliable. The results are displayed in Table D1 columns (I) & (II). The shares of exact freezes and cuts in weekly basic pay are not notably different from the shares found using basic (hourly) wages in a 0.005 range around zero (columns (III) & (IV), Table D1). If anything, basic wages appear less rigid than basic weekly pay, suggesting our approach of using a range around zero to define a freeze is a conservative choice. Since weekly basic pay changes can also occur when basic hours worked change, we repeat the same exercise for a sub-sample of job stayers who recorded an exactly zero year-to-year change in hours worked (around 20 percent of salaried job stayers). The results are displayed in Table D2 columns (I) & (II). The average spike at zero is 18 percent and the frequency of cuts is only 7 percent, suggesting that our adjustment in the main results of using a 0.005 range around zero

[4]This is no hypothetical issue. The authors know from personal employment experience that some employers in the Scottish hospitality sector frequently make mistakes when converting decimal values into hours and minutes worked.

provides a lower bound on the true proportion of basic wage freezes and an upper bound on the share of cuts.

What might have been the impact of the above sources of measurement error on the findings from previous studies of DNWR in Britain? To answer this question, we first compute the shares of *exact* freezes and cuts in job stayers' average earnings per hour, excluding overtime. This is the same measures focused on by Nickell and Quintini (2003) and Elsby, Shin, and Solon (2016) in their analyses on the extent of DNWR in the New Earnings Survey Panel Dataset (NESPD), 1975-2012, the precursor of the ASHE. We also ignore the distinction between salaried and hourly-paid workers to make this comparison. Table D3 (I)-(II) shows the results and confirms the previous payroll-based findings for Great Britain (Nickell and Quintini, 2003; Elsby, Shin, and Solon, 2016): nominal earnings per hour, excluding overtime, are reported as being cut quite frequently and the data show relatively few exact pay freezes, even during the recession and low inflation period following the 2007-08 global financial crisis. Specifically, column (I) shows that the share of exact year-to-year freezes in the data ranges from 3.1 percent in 2006-07 to 9.9 percent in 2012-13. For the 2006-12 period, when the two analyses overlap, our estimated average shares of *exact* zero changes in earnings per hour, excluding overtime, and cuts are within a range of one percentage point of what Elsby, Shin, and Solon (2016) found. The small deviations from their results are likely caused by a coding error. According to the published replication file, Elsby, Shin, and Solon used the 'same job' marker in the data to select job stayers, and then studied the difference between current wages and the last record of an individual, without conditioning that there was just one year between these records. This error in Elsby, Shin, and Solon marginally raises the measured frequency of nominal wage cuts and biases down the frequency of exactly zero changes throughout the Great Britain part of the study, typically within the range of one percentage point. We use a range of 0.5 log points around zero to address the measurement issues of hours worked described above. Table D3 (III)-(IV) shows freezes of between 7.4 percent (2007-08) and 20.6 percent (2009-10), substantially higher than the comparable estimates in Elsby, Shin, and Solon. This suggests that marginal reporting errors in the values of hours worked might have had a large impact on the conclusions reached by previous research.

TABLE D1: Nominal changes in weekly basic pay and basic wages, salaried job stayers.

Years	Weekly basic pay		Basic wages	
	Exact Freezes (%) (I)	*Exact* Cuts (%) (II)	Freezes (%) (III)	Cuts (%) (IV)
2006-07	10.3	9.0	10.0	11.3
2007-08	7.8	7.7	7.5	10.0
2008-09	15.6	10.2	15.0	11.1
2009-10	21.8	12.5	21.1	14.5
2010-11	21.3	10.8	21.4	11.9
2011-12	27.3	10.3	27.2	12.2
2012-13	25.1	10.1	24.3	12.3
2013-14	13.8	8.6	13.9	9.9
2014-15	12.5	9.1	13.0	10.1
2015-16	17.7	9.8	17.6	11.1
2016-17	13.9	9.3	14.4	10.8
2017-18	15.3	9.2	15.1	11.2
Average	16.9	9.7	16.7	11.4

Notes: Freezes and cuts show the percentage of job stayers with no change and negative change in the pay measure indicated. See Appendix Table E1 column (III) for annual sample sizes.

TABLE D2: Nominal changes in weekly basic pay with unchanged hours worked, salaried job stayers.

Years	Weekly basic pay, constant hours subsample		Basic wages	
	Exact Freezes (%) (I)	*Exact* Cuts (%) (II)	Freezes (%) (III)	Cuts (%) (IV)
2006-07	14.1	5.0	10.0	11.3
2007-08	8.3	5.0	7.5	10.0
2008-09	15.1	8.3	15.0	11.1
2009-10	22.9	11.3	21.1	14.5
2010-11	21.8	7.5	21.4	11.9
2011-12	29.6	8.0	27.2	12.2
2012-13	27.0	8.2	24.3	12.3
2013-14	14.1	6.4	13.9	9.9
2014-15	12.5	6.8	13.0	10.1
2015-16	18.2	7.6	17.6	11.1
2016-17	14.8	7.1	14.4	10.8
2017-18	16.5	7.0	15.1	11.2
Average	17.9	7.4	16.7	11.4

Notes: Freezes and cuts show the percentage of job stayers with no change and negative change in the pay measure indicated.

TABLE D3: Nominal changes in job-stayer average earnings per hour, excluding overtime.

Years	Earnings per hour, excl. overtime		Earnings per hour, excl. overtime	
	Exact Freezes (%) (I)	*Exact* Cuts (%) (II)	Freezes (%) (III)	Cuts (%) (IV)
2006-07	3.1	20.3	9.1	16.5
2007-08	3.2	17.9	7.4	14.4
2008-09	5.8	20.4	12.6	16.5
2009-10	8.7	24.3	17.7	18.8
2010-11	7.8	22.8	18.2	17.2
2011-12	9.7	23.4	20.6	17.0
2012-13	9.9	23.0	19.5	17.3
2013-14	6.5	20.4	13.8	16.1
2014-15	5.5	18.9	12.1	14.7
2015-16	6.6	20.7	14.5	16.1
2016-17	6.0	19.8	12.7	15.3
2017-18	6.4	18.8	13.8	14.4
Average	6.8	21.0	14.6	16.3

Notes: Freezes and cuts show the percentage of job stayers with year-to-year no change and a negative change in the pay measure indicated. See Appendix Table E1 column (I) for annual sample sizes.

Appendix E. Additional Tables and Figures

TABLE E1: Sample sizes of job-stayer wage change observations

Years	Total (I)	Hourly-paid (II)	Salaried (III)
2006-07	64,618	20,962	43,656
2007-08	64,558	21,494	43,064
2008-09	65,010	21,854	43,156
2009-10	82,258	29,140	53,118
2010-11	81,920	30,240	51,680
2011-12	81,973	30,603	51,370
2012-13	80,707	28,911	51,796
2013-14	83,362	28,672	54,690
2014-15	82,159	28,921	53,238
2015-16	78,952	27,578	51,374
2016-17	75,193	25,847	49,346
2017-18	74,086	25,865	48,221

Notes: Displays annual sample sizes of job stayers for the different sub-samples used.

TABLE E2: Nominal changes in basic wages, salaried job stayers

Years	Freezes (%) (I)	Cuts (%) (II)	LSW (%) (III)	Unemployment rate (%) (IV)	Inflation rate (%) (V)
2006-07	10.0	11.3	13.3	5.4	2.7
2007-08	7.5	10.0	13.2	5.4	2.9
2008-09	15.0	11.1	15.3	5.2	2.4
2009-10	21.1	14.5	14.1	7.6	3.7
2010-11	21.4	11.9	14.5	7.9	4.4
2011-12	27.2	12.2	14.2	7.8	3.0
2012-13	24.3	12.3	13.8	8.1	2.4
2013-14	13.9	9.9	16.7	7.8	1.8
2014-15	13.0	10.1	16.9	6.4	−0.2
2015-16	17.6	11.1	19.1	5.6	0.3
2016-17	14.4	10.8	18.9	4.9	2.7
2017-18	15.1	11.2	15.4	4.4	2.4
Average	16.7	11.4	15.5	6.4	2.4

Notes: Freezes and cuts show the percentage of job stayers with year-to-year no change and a negative change in basic wages. See Table E1 column (III) for annual sample sizes. 'LSW' refers to the Lebow, Saks, and Wilson (2003) statistic, as described in the main text, providing a measure for the asymmetry of the wage change distribution. Inflation is measured as the April-to-April log change in the UK Consumer Price Index (CPI). The unemployment rate refers to UK individuals aged 16 and over, seasonally adjusted and for the second April of each period, expressed as a percentage of the economically active population. Both series are from the Office for National Statistics.

TABLE E3: Nominal changes in basic wages, hourly-paid job stayers

Years	Freezes (%) (I)	Cuts (%) (II)	LSW (%) (III)	Unemployment rate (%) (IV)	Inflation rate (%) (V)
2006-07	14.6	4.7	12.4	5.4	2.7
2007-08	11.9	3.4	13.5	5.4	2.9
2008-09	18.1	5.5	12.9	5.2	2.4
2009-10	26.5	4.6	19.3	7.6	3.7
2010-11	28.0	3.9	12.8	7.9	4.4
2011-12	27.4	3.4	13.1	7.8	3.0
2012-13	25.0	6.3	11.2	8.1	2.4
2013-14	22.7	3.2	16.9	7.8	1.8
2014-15	16.8	3.1	13.2	6.4	−0.2
2015-16	19.5	5.7	22.8	5.6	0.3
2016-17	17.9	2.8	13.4	4.9	2.7
2017-18	19.8	2.7	14.1	4.4	2.4
Average	20.7	4.1	14.6	6.4	2.4

Notes: Freezes and cuts show the percentage of job stayers with year-to-year no change and a negative change in basic wages. See Table E1 columns (II) for annual sample sizes. 'LSW' refers to the Lebow, Saks, and Wilson (2003) statistic, as described in the main text, providing a measure for the asymmetry of the wage change distribution. Inflation is measured as the April-to-April log change in the UK Consumer Price Index (CPI). The unemployment rate refers to UK individuals aged 16 and over, seasonally adjusted and for the second April of each period, expressed as a percentage of the economically active population. Both series are from the Office for National Statistics.

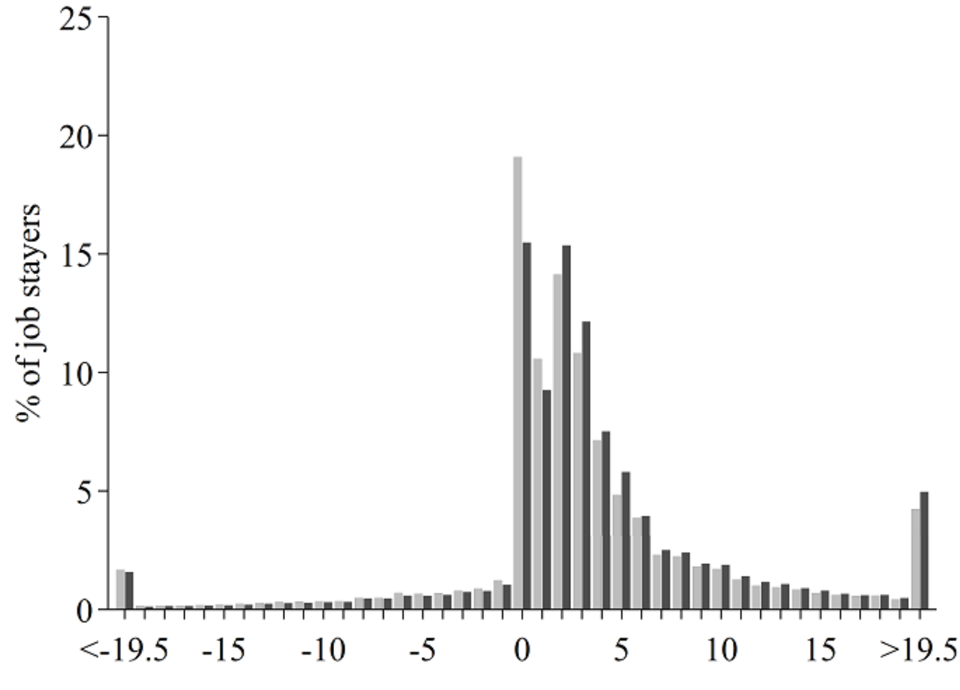

FIGURE E1: Frequency distribution of year-to-year changes in log basic wages, shrinking vs expanding firms, job stayers

Notes: See Figure 2. Pooled data for 2006-2018. **Light bars:** 277,001 job stayers whose firms experienced *negative* year-to-year employment growth over the same period. **Dark bars**: 259,236 job stayers whose firms experienced *positive* year-to-year employment growth over the same period.

Appendix F. Year-To-Year Adjustments in Earnings Per Hour

In the main text, we presented evidence on DNWR using basic wages. We have argued that basic wages are the relevant wage notion for macroeconomic models that include stickiness in nominal wages to generate fluctuations in unemployment. Basic wages exclude various kinds of extra pay: shift premium pay, overtime pay, commissions and incentive pay, and other pay (e.g., meal and travel allowances). In this Appendix, we analyse changes in average earnings per hour, excluding overtime, for further comparison with the existing literature where data availability did not permit an analysis of basic wages. This is the same wage measure as previously analysed by Nickell and Quintini (2003) and Elsby, Shin, and Solon (2016). For a discussion on the likely effects of measurement error in hours worked on the results presented in the previous studies, and how our proposed corrections lead to different estimates of wage freezes and cuts, see Appendix D.

How should changes in earnings per hour, excluding overtime, be interpreted? Changes in labour income caused by changes in working hours do no not affect a firm's marginal labour input cost per hour (except for the very first overtime hour). Similarly, changes in commissions and incentive pay or allowances are likely linked to a worker's output (e.g., sales) and labour input (e.g., business travel). Again, the marginal production costs a firm faces are likely unaffected by these extra pay components. However, from a worker's perspective, changes in earnings per hour might be perceived as income risk, depending on whether workers welcome reductions in long working times or less business travel (Devereux, 2001; Jardim, Solon, and Vigdor, 2019).

We begin by describing the year-to-year changes in job-stayer earnings per hour, excluding overtime. Figure F1 displays the results for hourly-paid (dark bars) and salaried (light bars) job stayers. The patterns in Figure F1 are similar to the ones seen in Figure 2 for basic wages. Nominal cuts are somewhat more common; the average share of hourly-paid job stayers receiving cuts is 11 percent, and for salaried job stayers this share is 17 percent. However, the distributions show strong signs of a censoring of earnings cuts. The spike at zero is substantial, though slightly lower than for basic wages. On average, 17 percent of hourly-paid and 14 percent of salaried job stayers have their earnings per hour, excluding overtime, frozen from year-to-year. The numbers imply constant quarterly wage-change probabilities of 0.36 and 0.39, respectively.[5] However, we caution against taking these probabilities at face-value as evidence against nominal rigidity, given the above described ambiguity in interpreting nominal changes in earnings per hour, excluding overtime.

Figure F2 displays time series evidence for the earnings per hour, excluding overtime, adjustments of job stayers comparable to Figure 3 for basic wages. The share of year-to-year freezes is strongly rising throughout the recession period for hourly-paid and salaried job stayers, peaking at 21.2 percent and 21.9 percent, respectively, in 2011-12. The Lebow, Saks, and Wilson (2003) statistic, indicating the missing mass of cuts, does not provide clear evidence for binding DNWR in earnings per hour, excluding overtime, for salaried workers. For hourly-paid workers, the missing mass falls from 12.2 percent to 9.1 percent from 2008 to 2009, but substantially increases to 14.1 percent for job stayers from 2009 to 2010.

[5] As before, we convert values as follows: $p_q = 1-(1-p_y)^{1/4}$, with p_q denoting the quarterly wage-change probability, and p_y the yearly wage-change probability.

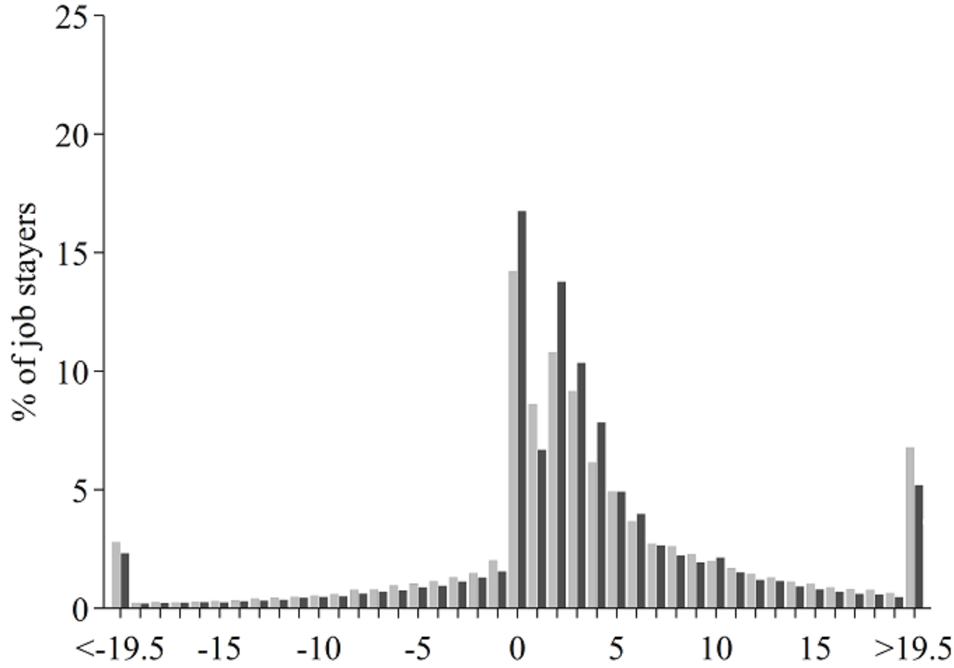

FIGURE F1: Frequency distribution of year-to-year changes in log earnings per hour, excluding overtime, job stayers

Notes: See Figure 2. **Light bars**: salaried. **Dark bars**: hourly-paid.

Finally, we repeat our two-step estimation from Section 5 to assess whether there is evidence that employers use extra pay components to circumvent constraints from internal wage structures. Table F1 shows that the earnings per hour, excluding overtime, of both incumbent employees and new hires, significantly decrease when the regional unemployment rate increases by one percentage point, 0.36 and 0.34 percent, respectively. As for basic wages (see Table 5), the difference in cyclicality is neither economically nor statistically significant (column III).

Our results for job-stayer adjustments in earnings per hour, excluding overtime, support the findings in the literature from studies of employer payroll records or payslips, recently summarised by Elsby and Solon (2019): cuts in average earnings per hour occur commonly, typically affecting more than 15 percent of job stayers. We view this as important evidence that workers face considerable earnings risk. Our results are also in line with findings for the US from Fallick, Villar, and Wascher (2020). They used data from the Employment Cost Index of the US Bureau of Labor Statistics, where the unit of analysis is not the worker, but the job. For each job, they observed the average nominal compensation per employee, including benefits such as pension payments. They reported an average proportion of year-to-year compensation freezes of 16 percent and nominal cuts of 15 percent. Because they cannot control for any changes to worker characteristics in the jobs, average nominal compensation changes can occur when the unobserved underlying skills of the employees in a job change. Consistent with that, they find much more variation in nominal compensation in jobs than we find here.

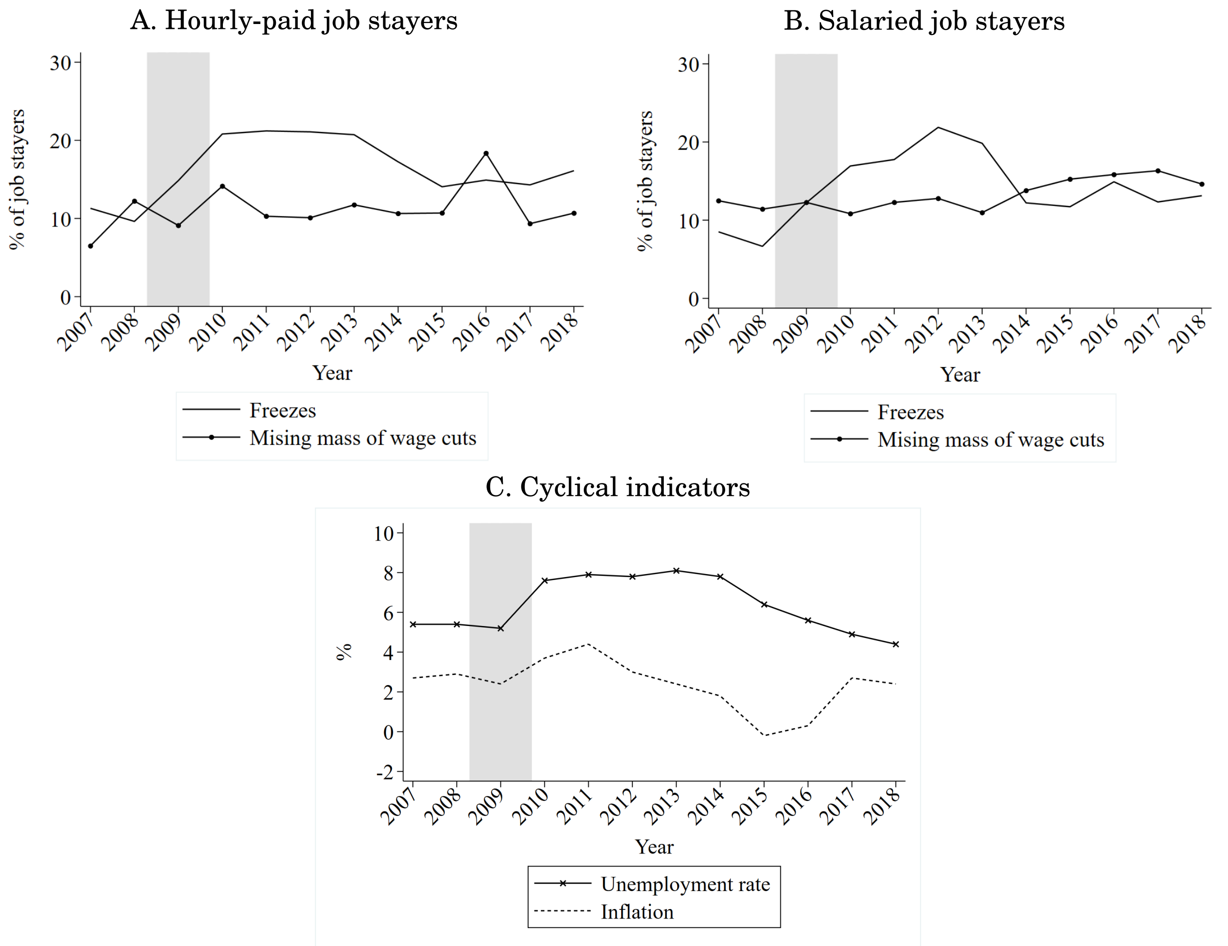

FIGURE F2: Quantitative indicators of DNWR, earnings per hour, excluding overtime

Notes: 'Freezes' show year-to-year changes in log earnings per hour, excluding overtime, for job stayers in the interval $[-0.005, 0.005]$. 'Missing mass of cuts' show the LSW statistic. See Appendix Tables E1 for sample sizes by year. Inflation is measured as the April-to-April log change in the UK Consumer Price Index (CPI). The unemployment rate refers to UK individuals aged 16 and over, seasonally adjusted and for the second April of each period, expressed as a percentage of the economically active population. Both series are from the Office for National Statistics.

TABLE F1: Estimated nominal wage responses of new hires and job stayers to regional unemployment and gross value added

Dependent variable / Cyclical indicator	Incumbents (≥ 1 year in job) (I)	New hires (< 1 year in job) (II)	Difference (Stayer-hire) (III)
Unemployment rate:			
1. Earnings per hour, excl. overtime	−0.0036	−0.0034	0.0001
	(0.0013)	(0.0017)	(0.0018)
Nominal GVA:			
2. Earnings per hour, excl. overtime	0.983	0.621	0.361
	(0.220)	(0.777)	(0.752)

Notes: Standard errors in parentheses robust to two-way clustering over year and NUTS1 region.
Sample sizes: first-step estimates for earnings per hour, excluding overtime, 1,079,602. Second-step estimates based on 121 year-region observations.
Column (I) shows estimates of the cyclical response of earnings per hour, excluding overtime, from the second-step regressions for job stayer wages. Column (II) shows the equivalent estimates for new hire earnings per hour, excluding overtime. Column (III) shows estimates from second-step regressions where the dependent variable is the difference between the incumbent and new hire coefficients from the first step, i.e., $\hat{\beta}^I_{rt} - \hat{\beta}^N_{rt}$.
Rows 1 shows semi-elasticity estimates with respect to the unemployment rate.
Rows 2 shows elasticity estimates with respect to gross value added.
Sources: Earnings per hour, excluding overtime, are from the Annual survey of Hours and Earnings. NUTS1 unemployment rates are from the Office for National Statistics (ONS) for April of each year, corresponding to the reporting period of the ASHE. Regional GVA are also from the ONS and correspond to calendar years, i.e., April-to-April changes in wages are regressed on the annual change in GVA over the years prior to each April.

Appendix G. Further Details on the Heterogeneity in Wage Freezes and Cuts across Workers and Firms

We study the conditional likelihood of year-to-year nominal wage cuts or freezes of job stayers. The following process describes whether a wage cut or freeze between periods $t-1$ and t is observed for job-stayer i:

$$y_{it} = \begin{cases} 1 & \text{if} \quad y^*_{it} > 0, \\ 0 & \text{if} \quad y^*_{it} \leq 0 \end{cases} \tag{7}$$

whereby the latent variable y^*_{it} for job stayer i is

$$y^*_{it} = \mathbf{x}'_{it}\boldsymbol{\beta} + \varepsilon_{it}, \qquad \varepsilon_{it} \sim N(0,1) \tag{8}$$

Here, y_{it} either represents basic wage freezes or cuts. $\mathbf{x}_{it}$ includes dummy variables indicating an employee's gender, whether their employer is in the private or non-private sector, whether their wage is set according to any form of collective agreement, whether the job is in the 'Wholesale & Retail Trade, Hotels & Restaurants sectors' industry (SIC G-H), and the worker's age band. We also include dummies in $\mathbf{x}_{it}$ for employer sizes: 'small' ($<$ 50 employees), 'medium' (50-249 employees), 'large' (500-4,999 employees), and 'very large' ($\geq 5,000$ employees). Additionally, we include dummy variables for the employee's position in the basic weekly pay distribution among all ASHE employees in the respective year, whereby 'low earnings' or 'high earnings' indicate that the employee earned below 2/3 or above 4/3 of median basic weekly wages that year, respectively, and 'medium earnings' refers to all remaining job stayers. All these dummy variables refer to period $t-1$ values. We also include indicator variables for the year-to-year growth in the number of employees at a job-stayer's firm, over the corresponding period (shrinking, expanding or constant), which proxy for the state of the firm.[6] For hourly-paid job stayers, we further include a dummy variable indicating whether the reported hourly basic wage is a multiple of ten pence ('round'). The results of estimating the probit models are displayed in Tables G1 and G2, with both slope coefficients and margins (probabilities) displayed, with the latter evaluated at the sample averages using the delta method.

[6] The employee counts of firms are provided in the ASHE dataset from the Inter-Departmental Business Register (IDBR), which is the official list of UK enterprises. We use the more common term 'firm' interchangeably with 'enterprise', which refers to a UK-specific administrative definition of an employer that could contain several local units or plants.

TABLE G1: Probit estimates: Incidence of year-to-year basic wage freezes for job stayers

	Salaried		Hourly-paid	
	Coefficient (I)	Margin (Pr.) (II)	Coefficient (III)	Margin (Pr.) (IV)
1. Female		0.162		0.190
2. Male	0.005	0.183	0.110^{***}	0.222
3. Non-private sector		0.176		0.234
4. Private sector	-0.023^{***}	0.170	-0.126^{***}	0.197
5. No coll. agree.		0.181		0.226
6. Coll. agreement	-0.063^{***}	0.165	-0.144^{***}	0.185
7. Age 15-29		0.108		0.156
8. Age 30-44	0.258^{***}	0.164	0.179^{***}	0.203
9. Age 45-64	0.408^{***}	0.204	0.265^{***}	0.228
10. Low earnings		0.185		0.181
11. Medium earnings	0.009	0.188	0.137^{***}	0.219
12. High earnings	-0.100^{***}	0.160	0.141^{***}	0.221
13. Small firms		0.267		0.342
14. Medium firms	-0.274^{***}	0.192	-0.294^{***}	0.242
15. Large firms	-0.409^{***}	0.157	-0.495^{***}	0.184
16. V. large firms	-0.445^{***}	0.149	-0.626^{***}	0.151
17. SIC G-H		0.179		0.169
18. Not SIC G-H	-0.025^{***}	0.172	0.199^{***}	0.224
19. Firm expanding		0.149		0.183
20. Firm const. size	0.122^{***}	0.178	0.071^{***}	0.203
21. Firm shrinking	0.186^{***}	0.196	0.176^{***}	0.234
22. 'Non-round'				0.192
23. 'Round' rate			0.193^{***}	0.249
N of job stayers	503,330		272,648	
Log-likelihood	-230,252		-134,779	

Notes: The dependent variable is an indicator that takes the value one if a job stayer experienced a basic wage freeze from year-to-year, and zero otherwise. Marginal effects are evaluated at the sample average using the delta method. Excludes all job stayers for 2007-2008 and 2017-2018 because firm size was incorrectly recorded by the ONS in the second years of those periods.

***, **, * indicate significance from zero of the model coefficients at the 0.1%, 1% and 5% levels, respectively, two-sided tests and robust standard errors.

TABLE G2: Probit estimates: Incidence of year-to-year basic wage cuts for job stayers

	Salaried		Hourly-paid	
	Coefficient (I)	Margin (Pr.) (II)	Coefficient (III)	Margin (Pr.) (IV)
1. Female		0.110		0.034
2. Male	0.042^{***}	0.118	0.161^{***}	0.048
3. Non-private sector		0.109		0.043
4. Private sector	0.046^{***}	0.118	-0.051^{***}	0.039
5. No coll. agree.		0.117		0.041
6. Coll. agreement	-0.033^{***}	0.111	-0.020^{*}	0.039
7. Age 15-29		0.090		0.036
8. Age 30-44	0.155^{***}	0.117	0.079^{***}	0.042
9. Age 45-64	0.160^{***}	0.118	0.055^{***}	0.040
10. Low earnings		0.166		0.046
11. Medium earnings	-0.206^{***}	0.119	-0.060^{***}	0.041
12. High earnings	-0.305^{***}	0.101	-0.249^{***}	0.027
13. Small firms		0.134		0.035
14. Medium firms	-0.106^{***}	0.112	0.036^{**}	0.038
15. Large firms	-0.129^{***}	0.108	0.005	0.035
16. V. large firms	-0.112^{***}	0.111	0.127^{***}	0.046
17. SIC G-H		0.109		0.023
18. Not SIC G-H	0.027^{***}	0.114	0.353^{***}	0.050
19. Firm expanding		0.113		0.035
20. Firm const. size	−0.002	0.113	0.027	0.037
21. Firm shrinking	0.009	0.115	0.149^{***}	0.048
22. 'Non-round'				0.041
23. 'Round' rate			-0.061^{***}	0.036
N of job stayers	503,330		272,648	
Log-likelihood	-178,677		-47,216	

Notes: The dependent variable is an indicator that takes the value one if a job stayer experienced a basic wage cut from year-to-year, and zero otherwise. Marginal effects are evaluated at the sample average using the delta method. Excludes all job stayers for 2007-2008 and 2017-2018 because firm size was incorrectly recorded by the ONS in the second years of those periods.

***, **, * indicate significance from zero of the model coefficients at the 0.1%, 1% and 5% levels, respectively, two-sided tests and robust standard errors.