Essays on Labour and Development Economics

Daniel Schaefer

A thesis presented for the degree of
Doctor of Philosophy

School of Economics
College of Humanities and Social Sciences
University of Edinburgh
2018

Lay Summary

This thesis presents three essays, each seeking to deepen our understanding of labour markets. The first essay studies the responses of real wages and hours worked of new hires to changes in the unemployment rate during the UK's Great Recession. The responses of these variables are interesting, because a firm's hiring decision depends to a large extent on the costs that it incurs if it hires a new employee. I find that real wages and hours worked of new hires significantly declined when the unemployment rate increased. This can potentially explain the relatively small decline of employment that was observed during the Great Recession: firms kept hiring because of falling costs and hours worked of new hires.

Another important determinant of a firm's hiring decision is its expectation about future labour market conditions. For example, if a firm expects that it will be relatively costly to find new employees in the next year, then this firm might already increase recruiting and worker-retaining efforts today. Usually it is assumed in economic models that firms form their expectations in a way which implies they predict future labour market conditions on average correctly. The second essay studies in how far the short-run labour market dynamics and long-run outcomes, for example, the long-run level of the unemployment rate, change if firms and unemployed workers have to use simple statistical methods to form forecasts, rather than "knowing" it. I show that the long-run outcomes are mostly unchanged, but the short-run dynamics of unemployment depend on the assumptions made about the formation of expectation.

The final essay computes the share of labour income in total income, for both the services and the goods sector, in a large cross-section of countries. The labour income shares of both sectors increase across countries with the level of development, measured by real output per person. Because no comparable data on these shares were available across countries, economists usually assumed that these income shares do not differ from their corresponding U.S. values. This assumption resulted in incorrectly measured productivity levels: the gaps in productivity levels across countries are larger than previously computed in other studies.

Abstract

This thesis presents three essays, each seeking to deepen our understanding of labour markets. The first essay studies the response of real wages and hours of new hires to the business cycle during the UK's Great Recession. The second essay analysis in how far the assumption of rational expectations in the Mortensen-Pissarides model is required for the economy to converge to an equilibrium. In particular, it asks if it is possible for economic agents to use simple linear forecast rules and still ensure convergence to the rational expectations equilibrium. The final essay seeks to determine whether labour income shares at the sectoral level are constant across countries, as is usually assumed in the literature, and whether this assumption quantitatively matters. Therefore, it takes the input-output structures across countries into account, and conducts a development accounting exercise.

Real wages and hours in the Great Recession: Evidence from firms and their entry-level jobs

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

Unemployment and econometric learning

I apply well-known results of the econometric learning literature to the Mortensen-Pissarides real business cycle model. Agents can always learn the unique rational expectations equilibrium (REE), for all possible well-defined sets of parameter values, by using the minimum-state-variable solution to the model and decreasing gain

learning. From this perspective, the assumption of rational expectations in the model could be seen as reasonable. But using a parametrisation with UK data, simulations show that the speed of convergence to the REE is slow. This type of learning dampens the cyclical response of unemployment to small structural shocks.

Measuring sectoral income shares: Accounting for input-output structures across countries

I use input-output tables to measure the labour income shares of the goods and the services sector for a large cross-section of mostly developed countries. I present two novel findings: sectoral labour income shares significantly increase with the level of development, and within-country differences between these income shares are uncorrelated with the level of development. These cross-country differences are not caused by variation in the input-output structure or final demand, but originate at the production-side of the economy. I measure sectoral total factor productivity using a development accounting framework to assess the quantitative importance of my findings. The goods sector of less developed countries is relatively less productive than the services sector; assuming that the values of the sectoral labour income shares across countries are identical to their corresponding U.S. values leads to an underestimation of productivity differences across countries. All findings are robust to different adjustments for the labour income of the self-employed.

Declaration

I, Daniel Schaefer, confirm that the work presented in this thesis is my own. Where the research was carried out alongside others, or where information has been derived from other sources, I confirm that this has been indicated in the thesis. This work has not been submitted for any other degree or professional qualification.

Daniel Schaefer

Edinburgh, 18 December 2017

Acknowledgements

I thank my academic supervisors, Andy Snell and Jan Grobovšek for their patience, support and advice, and I would express similar sentiments to many other academic and non-academic staff members in the School of Economics. I also thank Carl Singleton for productively working alongside me to produce parts of this thesis.

Contents

List of Figures

List of Tables

Introduction

The UK's Great Recession was the most severe economic contraction since the Second World War, yet employment declined less than in previous recessions. This resilience of UK employment has been attributed to flexible labour costs (Blundell et al., 2014), but economy-wide averages tell nothing about the responses of wages and hours *within* jobs, which is what determines a firm's employment decision in frictional labour markets (Pissarides, 2009). For example, if workers switch from high- to low-paying jobs during recessions, and vice versa during booms, then worker-level regressions will find very flexible and procyclical wages, even if wages within jobs are unchanged.

In the first essay I use the Annual Survey of Hours and Earnings (ASHE), a matched employer-employee dataset, to control for this cyclical job-switching. Only three other studies exist that control for job-switching when estimating the response of new hires' wages to the business cycle: Carneiro et al. (2012) and Martins et al. (2012), who use Portuguese data, and Stüber (2017) who uses German data. This essay offers multiple contributions: first, it extents the list of countries, for which results of new hires' wage flexibility exist, to the UK. This is important, because the UK is generally considered to have the most flexible labour markets among European countries, and thus results for countries with relatively strong unions and labour protection laws, such as Portugal and Germany, might not apply to the UK. The second contribution is that I provide robust job-level measures of the responses of real wages and hours worked for both, new hires and job stayers. In particular, I compute median real wages and hours of new hires and job stayers in each job (firm-occupation pair) and year, which extends the framework of Martins et al. (2012) by allowing for a direct comparison with job stayers' wages and hours in the same firm. Third, this essay provides novel evidence that job stayers' hours do not respond to the Great Recession, but new hires' hours are significantly reduced in large firms, mostly due to a shift from full-time to part-time hiring.

I estimate the response of these job-level variables to the aggregate unemployment rate, controlling for the changing composition of jobs over the business cycle by including job-fixed effects. However, I follow Martins et al. (2012) and select a particular sample of jobs for which I repeatedly observe hiring over the business cycle to ensure that my estimates are not affected by endogenous sample selection: for example, if jobs with more rigid hiring wages stopped hiring during recessions, then this would lead to an overestimation of the flexibility of hiring wages. In other words, I forego representativeness for certainty of what I'm measuring. I show that firms were able to respond to the Great Recession with substantial real wage cuts and by recruiting more part-time workers. A one percentage point increase in the unemployment rate led to an average decline in real hourly wages of 2.8 per cent for new hires and 2.6 per cent for job stayers. Hours of new hires in entry-level jobs in large firms were also substantially procyclical, while job-stayer hours were nearly constant. This substantial flexibility for new hires could explain the relatively high job-finding rate during the UK's Great Recession (Elsby and Smith, 2010). My findings suggest that models assuming rigid labour costs of new hires are not helpful for understanding the behaviour of unemployment over the business cycle.

The second essay studies the Mortensen-Pissarides business cycle model. Because of the absence of a centralised market for labour, the assumption implicit in this model that firms and unemployed workers "know" the fundamentals of their economy, such as the cost of posting vacancies or the equilibrium ratio of vacancies-to-unemployed, seems relatively strong. Therefore, I ask how the dynamics and equilibrium properties of the Mortensen-Pissarides model change when firms and unemployed workers have to use simple linear forecast rules, which are updated every period, to form expectations about future values of relevant variables. In other words, agents are assumed to "learn" the underlying parameters of the economy rather than to "know" them.

I apply well-known results of the econometric learning literature (Evans and Honkapohja, 2001), and show that agents can always learn the unique rational expectations equilibrium (REE), for all possible well-defined sets of parameter values. This means that the equilibrium quantities and ratios of the Mortensen-Pissarides model are unaffected by the behavioural assumption that agents are learning. In parallel work, Di Pace et al. (2016) arrive at similar results, however, in contrast to their work I use a simpler version of the Mortensen-Pissarides model, which enables me to derive analytical results instead of having to rely on numerical simulations. Therefore, this essay contributes to the literature on search in macroeconomic models of the labour market by showing that the standard assumption of rational expectations in these models seems to be a reasonable assumption with respect to the implied equilibrium quantities and ratios.

Additionally, this essay shows, by using a simulation which is parametrised with UK data, that the speed of convergence to the REE is slow. This means that, although the equilibrium will be reached asymptotically, aggregate variables such as the unemployment rate, could be persistently some distance away from their equilibrium values. Therefore, the rational expectations model of unemployment fluctuations could in fact be a poor approximation to an economy in which agents more realistically learn as econometricians, especially in the presence of frequent structural or permanent shocks. Moreover, this type of adaptive learning dampens the cyclical response of unemployment to small structural shocks, by making wages adjust more gradually after shocks, instead of jumping to their respective steady state values as under rational expectations.

The final essay measures the labour income shares of the goods sector and services sector for a large cross-section of mostly developed countries. The reason for this is that the literature on development economics and economic growth usually makes the assumption that these sectoral labour income shares are identical to the ones measured in the United States by Valentinyi and Herrendorf (2008). This assumption is not based on theoretical arguments or direct empirical evidence, but rather a necessity because of a lack of comparable data on sectoral income shares across countries. The only indirect evidence so far are the results provided by Gollin (2002), who found that the *aggregate* labour income shares are not correlated with the level of income.

The aim of this last essay is simple: to understand whether the assumption of constant sectoral labour income shares across countries holds. Therefore, I use input-output tables to measure the labour income shares of the goods and the services sector for a large cross-section of mostly developed countries, following the method of Valentinyi and Herrendorf (2008). I present two novel findings: sectoral labour income shares significantly increase with the level of development, and within-country differences between these income shares are uncorrelated with the level of development. The first finding does seem to contradict the results of Gollin (2002), which could be due to the absence of least developed countries from the sample of countries which I analyse.

This essay then decomposes the measured labour income shares to provide a better understanding of the causes of the observed differences. I find that the cross-country differences are not caused by variation in the input-output structure or final demand, but originate at the production-side of the economy. Consequently, value-added labour income shares provide a good, and less data-demanding, approximation of the sectoral labour income shares.

Finally, I measure sectoral total factor productivity using a development accounting framework, similar to Herrendorf and Valentinyi (2012), to assess the quantitative importance of my findings. The goods sector of less developed countries is relatively less productive than the services sector. Assuming that the values of the sectoral labour income shares across countries are identical to their corresponding U.S. values leads to an underestimation of productivity differences across countries. All findings are robust to different adjustments for the labour income of the self-employed. Therefore, my findings suggest that future research should not use the U.S. labour income shares in cross-country studies, but rather use the shares provided in this essay, or the aggregate labour income shares as approximation.

Chapter 1

Real wages and hours in the Great Recession: Evidence from firms and their entry-level jobs

Note: This chapter has also been published as CESifo Working Paper No. 6766, and was co-authored with Carl Singleton, who is a Post-doctoral Fellow at the University of Edinburgh, School of Economics; e-mail: carl.singleton@ed.ac.uk. Carl has agreed that this essay represents in the majority my work, and that it can appear within this thesis. In addition to those already acknowledged, I am especially grateful to Steven Dieterle, Mike Elsby, Jonathan Thomas, and Ludo Visschers for their advice and comments. This work was presented at the 2017 Aarhus Conference on Markets with Search Frictions, the 8th ifo Conference on Macroeconomics and Survey Data, the 42nd Simposio of the Spanish Economic Association, and the internal seminar of the School of Economics, University of Edinburgh. The data used in this chapter are accessible from the UK Data Service, having been collected by the Office for National Statistics (ONS). Neither the collectors of the data nor the Data Service bear any responsibility for the analysis and discussion of results in this chapter.

1.1 Introduction

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

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

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

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

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

Since Solon et al. (1994) it has been recognised that the measured business cycle response of the average real wage typically underestimates the true response because of composition bias: the share in total hours worked of low-wage workers decreases during recessions, inducing a countercyclical bias. Several studies have used longitudinal data to address this bias (see the survey by Abraham and Haltiwanger, 1995). Most recently, Elsby et al. (2016) find that real wages are procyclical for UK employees working in the same job for at least one year, with an especially large response to the Great Recession. Pissarides (2009) however argues that what matters for a firm’s hiring decision, and thus vacancy creation, are wages in new worker-firm matches. There exists some evidence that the wages of workers who change employers respond to the business cycle (see Bils, 1985; Shin, 1994; Devereux, 2001, and Gertler et al., 2016 for US evidence). Devereux and Hart (2006) and Hart and Roberts (2011) find that the cyclicality of wages for British job changers significantly exceeds that of stayers.

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

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

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

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

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

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

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

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

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

$$w_{jt} = \alpha_j + \beta_t + x'_{jt}\delta + \varepsilon_{jt} \ , \tag{1.1}$$

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

The parameter estimates $\hat{\beta}_t$ from (1.1) are a series of period-means of log wages, regression adjusted for changes in the composition of jobs in the sample. In the second step, we relate this series to the Great Recession by regressing it on the unemployment rate U_t:

$$\hat{\beta}_t = c_0 + c_1 t + \gamma U_t + e_t \ . \tag{1.2}$$

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

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

$$w_{qk} = \alpha_q + \beta_{T(qk)} + \lambda k + x'_{qk}\delta + \varepsilon_{qk} \ , \tag{1.3}$$

where $T(qk)$ is a function indicating that a job q is observed in period t, and x'_{qk} contains time-varying characteristics: age squared and cubed, tenure in the firm

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

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

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

squared and cubed, firm size and its square, and the share of employees covered by a collective agreement. Linear terms for age and tenure are omitted as these would be collinear with k, and the average effect of these variables is controlled for by the estimated linear trend $\hat{\lambda}$. Although rewriting and estimating (1.3) in first differences over k is a more intuitive representation of how we estimate job-stayer wage cyclicality, we proceed with the equation in levels to obtain more directly comparable estimates of $\hat{\beta}_t$. The second step regressions for job-stayer and hiring wages are identical.

To measure the response of hours worked we estimate the same two-step models, replacing the dependent variables in (1.1) and (1.3).

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

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

The dataset can be viewed as a panel of employees without attrition, forming an approximate one per cent random and representative sample of UK employees in every year.[2] Particularly valuable for our analysis are the longitudinal identifiers for individuals (1997-2016) and enterprises (2003-2016). We use the terms “firm” and “enterprise” synonymously. The latter in this case is a specific administrative definition of UK employers, which could contain several local units (or plants). We believe this is the appropriate level to study firm- or job-level wages, because in most organisations pay-setting practices are determined at the enterprise level.

In another paper we use a combination of the exact number of employees, 4-digit industry classification and information on the legal status of an enterprise to define the boundaries of larger firms within each year back to 1996 (Schaefer and Singleton, 2017). But, here we want to link firms longitudinally before and after 2003. Therefore we impute values of the enterprise identifier backwards from 2003, using consecutive observations of individuals who have not changed jobs between years, as well as

[2]The two main reasons why an individual might not be observed in some year are: either being truly non-employed, or having changed employer between January and April. Since the survey questionnaires are in most cases sent in April to the employer’s registered address from January PAYE records, workers who switch employers during these months are undersampled.

using employment start dates. We then use the available within-year employer local unit identifiers to impute more enterprise values. For further information on how we construct this employer-employee panel from ASHE cross-sections, and other adjustments made to the data and the sample selection, see Appendix A.1.

Our analysis focuses on two main variables: basic weekly paid hours and the hourly wage rate, which equals the ratio of gross weekly earnings to the former, all excluding overtime. We refer to these simply as hours and wages. Monetary values are deflated using the Consumer Price Index (CPI).[3] For comparability with other studies, we include some statistics about nominal wage changes in Appendix A.5. We consider working-age employees (aged 16-64) in the private sector, who have non-missing records of earnings and hours. We include only the main job observation of an individual, which must not be at trainee or apprentice level, and not have incurred a loss of pay in the reference period for whatever reason. To avoid some spurious hourly wage rates we only keep observations with 1-100 basic paid weekly hours. Since the data are not top-coded, we drop the highest one per cent of weekly or hourly earners.

Our main indicator variable for the Great Recession is the working-age unemployment rate: the number of people unemployed divided by the economically active population.[4] To correspond with the timing of the ASHE, we use average values over the previous four quarters for all price series and business cycle indicators. For example, an estimate of the 2009 wage for new hires is compared with the average unemployment rate over the preceding twelve months, when those hires would have been made. We use the unemployment rate for comparability with the wider literature. In Section 1.4 we discuss the robustness of our main results to this choice.

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

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

[3]For robustness we also compute results using the Retail Price Index (RPI). All prices were obtained from UK National Statistics, accessed 24/04/2017.

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

in two different firms is treated as two separate jobs. We define a new hire as any employee with less than one year of tenure with a firm.

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

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

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

The entry-level jobs sample consists of 309 firms hiring into 391 jobs (Table 1.1). Our sample is unbalanced since some jobs are not observed in all years during 1998-2016. As Martins et al. (2012) note, the most important consideration is that the number of entry-level jobs should not vary systematically over the business cycle, as this would result in endogenous sample selection. The contemporaneous correlation of the number of entry-level jobs in our sample and the unemployment rate is insignificant (p-value: 0.49), and no other cyclical patterns are evident in Table 1.1 column (1). The median number of new hires per entry-level job is seven over the sample period.

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

TABLE 1.1: Number of new hires, entry-level jobs, and consistent-hiring-firms by year

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

Notes.- age 16-64, private sector only. Source of the unemployment rate series is discussed in Section 1.2.

observations of occupation-firm pairs, totalling 158,194 job stayers. The selected jobs represent on average nearly 90 per cent of all job stayers in the CH-firms sample over the whole period.

New hires in the CH-firms sample are younger, more likely to be female, and less likely to work full-time than job stayers (columns (1) and (2), Table 1.2). The wages and basic hours of new hires are lower than for job stayers. The same statements hold for the entire ASHE, (columns (3) and (4), Table 1.2), though the difference between the hours worked by all new hires and job stayers is considerably smaller: almost two-thirds of hires into entry-level jobs are part-time. When compared with the whole economy, the lower average age, real wages, and basic hours in the CH-firms sample can be explained by differences in industry and occupation composition. Over two-thirds of new hires are made by firms in the “accommodation and restaurant” and the “industrial cleaning and labour recruitment” industry. Similarly, the largest

TABLE 1.2: Descriptive statistics for employees: comparison of the consistent-hiring-firms sample and the whole ASHE (all firms and jobs), 1998-2016

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

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

shares are employed as service or sales workers (see Tables A.4-A.5 for the complete industries and occupation breakdowns). Firms in these industries account for approximately a third of all employees in the private sector (Jäger, 2016). This is also reflected in larger firms dominating the CH-firms sample. These large firms have average annual growth in the number of their employees of around four per cent, while the average for all firms is around eight per cent. Although the observable characteristics of new hires in entry-level jobs exhibit secular trends during 1998-2016, we do not see any notable cyclical patterns (Appendix Figure A.1).[5]

Since our subsequent analysis takes place at the job level, we compute median wages and hours within jobs each year (our results do not change significantly when we use average wages and hours instead). Figure 1.1 shows the distribution of real wages of new hires after subtracting their respective entry-level job median wages.

Although we can expect some dispersion around these median wages, the robustness of our econometric approach and the meaningful interpretation of any results to some extent depends on us capturing "typical" hiring wages. Some dispersion around the median hiring wages, indicated by zero, is visible in Figure 1.1. More than 50 per cent of hiring wages lie within a range of five log points, and almost 90 per cent within 10 log points around the job-specific median. The dispersion around the typical

[5] Appendix Figure A.1 shows that the share of men among new hires increases steadily by around ten per cent from 1998 to 2016, while the share of full-time employees decreases. As a consequence of our recursive sample construction prior to 2003, the average age of hires decreases by over five years from 1998 to 2003. Including controls for the average age of hires within a job in the following analysis does not change our results.

FIGURE 1.1: Distribution of differences between log real wages of new hires and their median values within entry-level jobs, 1998-2016

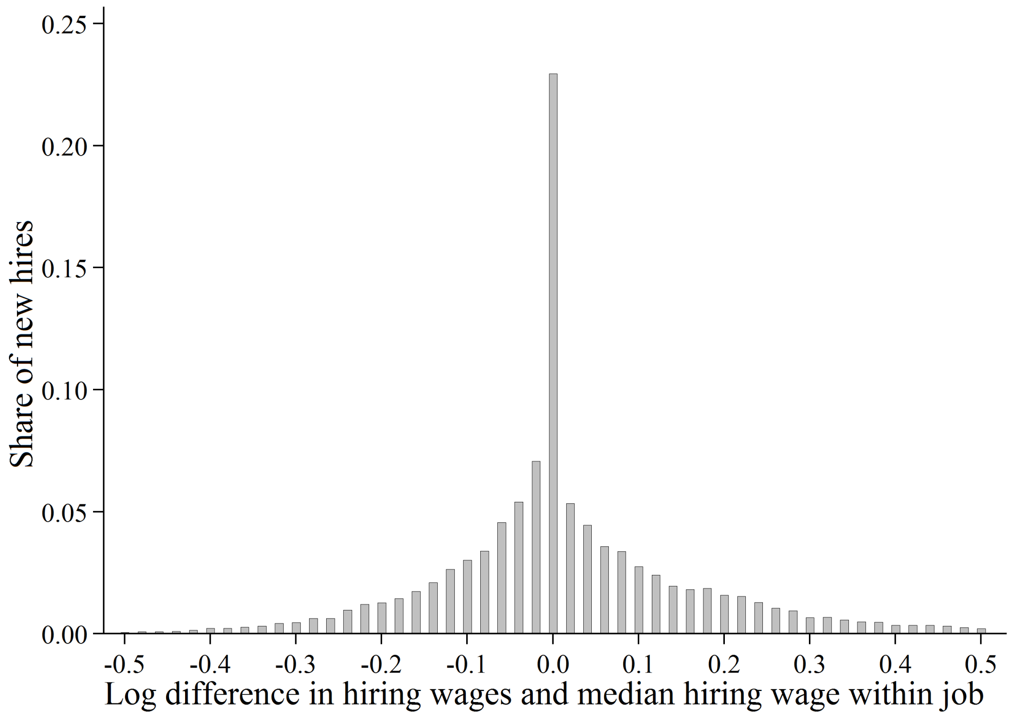

Notes.- within each entry-level job and year, the median hiring wage is subtracted from all hiring wages in that job and the resulting log differences are collected in bins with a width of two log points.

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

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

Figure 1.2 shows the estimated time series $\hat{\beta}_t$ for new hires and job stayers in CH-firms from regression models (1.1) and (1.3) over the period 1998-2016. The short-dashed lines indicate 95 per cent confidence intervals for the point estimates of $\hat{\beta}_t$ for new hires, using standard errors robust to serial correlation at the job-level. The confidence intervals for job stayers (not shown) are very narrow and lie within the intervals for new hires, except in 2015-16 for real wages and 2011-16 for hours worked. All series are normalised to zero in 2003 and series-specific linear trends

FIGURE 1.2: Estimated period-fixed effects for log real wages and log hours worked, including 95% confidence intervals for new hires, 1998-2016.

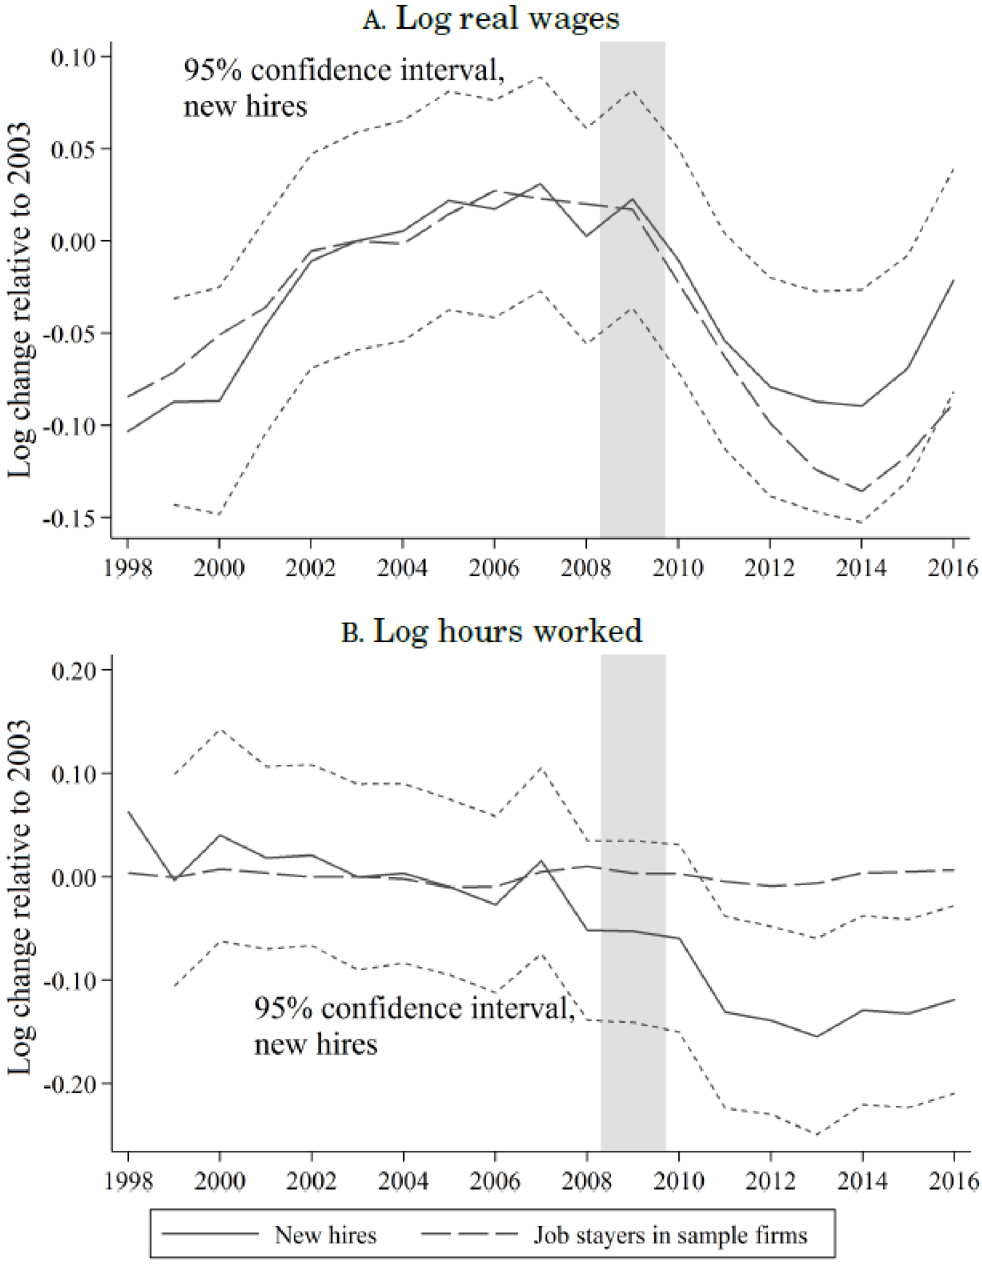

Notes.- the 95% confidence interval for job stayers (not shown) is very narrow and lies within the interval for new hires except in 2015-16 for real wages and 2011-16 for hours worked. Standard errors are robust to clustering at the firm-level. Excluded reference category in first-step regression (1.1) is 1998 for new hires, and regression (1.3) excludes 1998 & 2016. Series-specific linear trends removed from panel A. Series normalised to zero in 2003. Shaded area marks official UK recession dates. “New hires" are for wages in entry-level jobs where employees have less than twelve months of tenure. “Job stayers in sample firms" are for jobs and employees who have tenure greater than twelve months, and only for firms which are ever represented in the CH-firms sample.

have been removed from the series for wages for comparability. We estimate (1.1) with the unbalanced baseline panel of jobs described in Table 1.1, using period and job dummies. Thus, these time series should be interpreted as composition-adjusted real wages and hours of new hires and job stayers. In Appendix Figure A.5 we show the series for real wages without removing their respective linear trends.

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

We measure the response of real wages and hours to the Great Recession by estimating the second-step regression (1.2) using least squares. As recommended by Solon et al. (2015), we do not use weighted least squares (WLS) in our baseline regressions, because the least squares residuals do not display significant evidence of heteroskedasticity. Nevertheless, we later compare estimates from our baseline to those obtained using two different WLS estimators: (1) weights equal to the number of new hires; (2) weights equal to the number of entry-level jobs. Estimates from the first WLS estimator suffer from endogenous sample selection, because hiring volume is likely to depend on the cyclical response of wages and hours. However, the sign of the induced bias is informative about the bias in other studies using worker-level data. The second estimator is the WLS procedure applied by Martins et al. (2012), which accounts for the varying sample sizes of entry-level jobs over time.

The first row of Table 1.3 displays the main (or baseline) results, measuring the semi-elasticity with respect to a one percentage point (p.p.) increase in the unemployment rate: real hourly wages of new hires and job stayers decrease by 2.8 and 2.6 per cent if the unemployment rate increases by one p.p.. These estimates are significantly different from no response, but do not significantly differ from one another: Appendix Table A.3 shows the coefficient estimates when we regress the difference between new hires and job stayers in the estimated series of $\hat{\beta}_t$ on the

[6] Appendix Tables A.6-A.7 display the underlying values of all series in Figure 1.2.

unemployment rate, using our two-step approach. In column (3) we find that hiring hours respond by around 1.5 per cent, compared to only 0.2 per cent for job stayers in column (4). The significant decline in average hours per worker during the Great Recession has been discussed before (Blundell et al., 2014; Pessoa and Van Reenen, 2014; Borowczyk-Martins and Lalé, 2017), though not at the job level. Blundell et al. (2008) find that UK workers adjust hours worked in response to welfare reforms usually by changing firms, and this is particularly true for larger firms and in the services industry. To the best of our knowledge, the relatively greater and large response of hiring hours has not been documented previously.

TABLE 1.3: Estimated semi-elasticity of real wages and hours with respect to the unemployment rate, 1998-2016

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

Notes.- second-step regression results, estimates $\hat{\gamma}$: responses of the period-fixed effects $\hat{\beta}_t$ to the unemployment rate; regression specifications as in (1.1)-(1.3). First row refers to the main/baseline estimates. Second row includes an additional time-varying control for the share of full-time workers in a job. Third row changes the selection criteria for entry-level jobs, such that they have to be fulfilled in at least a quarter of years when the firm is observed, instead of a half, and those firms have to be observed for at least five years. Fourth row includes all job observations which hire in at least two years, or with at least two consecutive years of observations for job stayers. Fifth row uses WLS in the first-step, with weights proportional to the number of new hires or stayers in each job.

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

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

We also find that real weekly earnings (excl. overtime) of new hires decline by 4.7 per cent if the unemployment rate increases by one p.p., while job-stayer earnings decline by 2.9 per cent. Because the covariance between wages and hours is positive, these estimates exceed the sum of the corresponding values in the first row of Table 1.3.

We re-estimate regression (1.1), but in addition control for changes in the share of full-time workers within jobs. The second row of Table 1.3 shows that for hiring hours the semi-elasticity estimate falls to 0.7: over half of the recessionary decrease in hiring hours can be attributed to a shift from full- to part-time hiring. However, the response of real wages to the unemployment rate does not differ significantly between full- and part-time hires and job stayers.

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

As previously explained, the approach of Martins et al. (2012) has the advantage that we can be confident that what we measure is the response of real wages and hours of new hires in entry-level jobs. The potential disadvantage is that this response may not be representative of the whole economy. When we additionally include all other jobs in the ASHE for which we observe hiring in at least two periods (and similarly for job stayers), then the estimated time series of period-fixed effects from the first step resembles the series from our baseline sample (Appendix Figures A.3 and Figures A.4). The exception is hiring hours, which are less cyclical at the job level in the whole economy. This sample contains over twice as many hires and six times as many job stayers as our baseline, and the fourth row in Table 1.3 shows comparable estimates to our baseline. Real wages of all new hires and job stayers in this sample decrease by 2.5 and 2.9 per cent, respectively, for each p.p. increase in the unemployment rate. However the response of hiring hours is less pronounced in this larger but more representative sample, though still significant at the five per cent level. Thus, jobs with less flexible hiring wages and hours were more likely to stop hiring altogether during the Great Recession, inducing a countercyclical selection bias. In contrast, the selection bias for real wages of job stayers seems to be procyclical.

The final row of Table 1.3 shows values for $\hat{\gamma}$ when we estimate the first step using WLS, with weights proportional to the number of employees in each job. In this case the semi-elasticity estimate for hiring wages of 2.2 is the smallest of all the specifications described here. Therefore, weighting jobs by the number of hires induces a countercyclical bias in the estimated response of real wages to business cycle conditions: relatively more hires are made in jobs with relatively rigid hiring wages. Jobs which hired relatively more employees than others during the downturn also decreased the hours worked per hire more: the response to a one p.p. increase in the unemployment rate for entry-level hiring hours is more than 85 per cent larger than without endogenous weighting. Therefore, weighting by hiring volume overestimates the responsiveness of hiring hours. Not surprisingly, the firms who hire relatively more during recessions are also able to move their workforces towards shorter hours and part-time working. When Stüber (2017) similarly weights his job-level regression, German real hiring "wages" become more procyclical, seemingly contradicting our findings. But the results here could explain why we reach an opposite conclusion on the direction of bias induced by the endogenous hiring volume of jobs, which is also more in line with the hypotheses by Gertler and Trigari (2009) and Martins et al. (2012): the German data offers information on annual earnings and the number of days worked, and so Stüber's measures of real wages are better understood as average daily earnings. As our estimates show, the procyclical bias in hiring hours is large, and exceeds the countercyclical bias in hiring wages. When combined, this can cause a procyclical bias in earnings.

Finally, we make the sample selection criteria more exclusive, by increasing the required minimum number of hires in a job per year for it to be included as an entry-level job. The estimated semi-elasticity of hiring hours significantly increases in the number of minimum hires (see column (3) of Appendix Table A.2), nearly doubling when we require at least 10 hires per job and year. The estimated semi-elasticity of real hiring wages slightly increases in absolute terms, peaking at 2.9 when we raise the minimum number of hires to 7 per job and year. Varying the minimum number of hires generally affects the measured response to the Great Recession of both wages and hours. However, our main finding is unchanged: real wages of new hires are marginally more responsive to the unemployment rate than wages for job stayers in the same firms. Similarly, hiring hours always respond more strongly to the Great Recession than job-stayer hours.

1.4 Robustness and further discussion

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

1.4.1 Using other specifications of the regression model

Table 1.4 displays results from varying the specification of the second-step regression (1.2), while the first step remains unchanged. The main baseline results are repeated in the first row. The second row shows that when we include a quadratic time trend, wages decline marginally less when the unemployment rate increases, but the hours responses are approximately unchanged. We prefer to only include a linear trend because of the small number of periods in our dataset.

TABLE 1.4: Estimated semi-elasticity of real wages and hours with respect to the unemployment rate, 1998-2016: varying the specification of the second-step regression

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

Notes.- second-step regression results of estimated period effects on unemployment rate, $\hat{\gamma}$. First row is identical to Table 1.3, included here for comparison. Second row shows estimates when the second-step includes an additional quadratic time trend term. Third row estimates (1.2) in first differences, so measures the response of the log change in wages to a one percentage point increase in the change in unemployment. Fourth row applies WLS in the second step, with weights in proportion to the number of jobs observed per year.
Newey-West standard error estimates robust to first-order serial correlation in parentheses.
*** Statistically significant at the 1% level; ** at the 5% level, two-sided tests.

We also re-estimate (1.2) in first differences, to address potentially spurious estimates if wages, hours, or the unemployment rate are integrated. As in our baseline results, the real wage growth of new hires and job stayers does not respond significantly differently to changes in the unemployment rate: if the change in the unemployment rate increases by one p.p., then the growth of hiring wages decreases by 1.6 per cent, and for job-stayers' wages by 1.8 per cent. This is comparable to the finding by Devereux and Hart (2006) of 1.7 per cent for all job stayers in the UK during 1975-2001.

For comparability with Martins et al. (2012), we re-estimate (1.2) using WLS, with weights proportional to the number of jobs per period in the first step. The resulting estimates in the final row of Table 1.4 are qualitatively unchanged from the baseline. However real wages and hours are slightly less cyclical. Overall, our results that both the real wages of hires in entry-level jobs and of job stayers declined in response to the Great Recession are robust to the specification of the second-step regression. The finding that hiring hours declined more than for job stayers is also robust, except for the first-differenced version of (1.2), which indicates that the decrease in hiring hours is better understood as a medium-run and persistent development since 2008. In Appendix A.2 we discuss the results of further robustness checks, which also do not affect our confidence in the main results.

1.4.2 Using labour productivity as the business cycle indicator

As an alternative indicator for the Great Recession we consider labour productivity, measured by log real gross value added per hour.[7] Measures of labour productivity are particularly relevant for a firm's hiring decisions. As Haefke et al. (2013) explain, the estimated response to this measure has an intuitive interpretation in standard search and matching models of the labour market: if real wages are perfectly rigid, then they should not respond to labour productivity, while a one-to-one response indicates fully flexible wages.[8] Table 1.5 shows the estimated elasticity when we use labour productivity instead of the unemployment rate in regression (1.2). In the first row we use aggregate labour productivity as the business cycle indicator. The estimates are significantly smaller than one, but positive. Hiring hours also respond significantly, though less than real wages.

Because over 90 per cent of jobs in our baseline sample belong to the services industry, and the response of labour productivity to the Great Recession was not the same across sectors, we also use labour productivity of the services sector as the cyclical indicator. The second row of Table 1.5 shows a higher estimated elasticity of real wages and hours worked with respect to services sector labour productivity. Real wages of new hires and job stayers significantly decrease by 0.9 and 1.0 per cent when aggregate labour productivity decreases by one per cent. The difference between these values is insignificant. The estimates for job stayers and new hires do not significantly differ from one at the 99 per cent and 95 per cent confidence level

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

[8] Appendix Figure A.3 shows the time series of each business cycle indicator used and Appendix Table A.8 shows the underlying values.

TABLE 1.5: Estimated elasticity of real wages and hours with respect to labour productivity, 1998-2016

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

Notes.- second-step regressions of estimated period effects on alternative indicator of the business cycle. First-step estimated according to (1.1) and (1.3). "Labour productivity (I)" uses the log of real whole economy gross value added (GVA) per hour: ONS series LZVB. "Labour productivity (II)" uses the log of real gross value added (GVA) per hour in Services (sectors G-U): ONS series DJP9. We adjust both series by multiplying by the ratio of CPI to Producer Price Index of the services sector.
Newey-West standard error estimates robust to first-order serial correlation in parentheses.
*** Statistically significant at the 1% level; ** at the 5% level, two-sided tests.

respectively: suggestively, real wages in the UK were perfectly flexible to aggregate labour productivity over the sample period.

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

1.4.3 The role of the National Minimum Wage

Our results suggest that the real wages of new hires are just slightly more responsive to business cycle conditions than for job stayers. One potential explanation for this finding is the presence of a wage floor. This could constrain firms in how far they can reduce hiring wages. In 1999 such a floor was introduced in the UK in the form of the National Minimum Wage (NMW), with both adult and youth rates applying nationwide. These are usually uprated on an annual basis.[9] Collectively bargained wages can also limit a firm's flexibility in setting hiring wages. However, at the onset

[9] Source: https://www.gov.uk/national-minimum-wage-rates; accessed 01/07/2017.

of the Great Recession, only six per cent of new hires in our sample were covered by a national or industry-level collective agreement (affecting working conditions, not necessarily pay). Therefore we consider the NMW to be the more likely limit on the responsiveness of hiring wages.

Figure 1.3 displays the real NMW rate that applied to workers aged 21 and older, along with the 10th percentile, 25th percentile, and median real wages of new hires within entry-level jobs for each year.[10] These hiring wages are not adjusted for changes in sample composition and include only workers aged 22-64. Between 2006 and 2015 new hires at the 10th percentile of the wage distribution were paid the legal minimum, i.e. the real value of the adult rate. In 2016, the 10th percentile of new hires increased more than the adult rate, which followed the introduction of a higher NMW rate for workers aged 25 and over. We also observe a narrowing of the gaps between the minimum wage and both the 25th percentile and median of hiring wages over the sample period. In other words, the domain of the distribution of real hiring wages at the job level, for employees aged 22-64, became more restricted from below at the level of the real NMW adult rate during the recent downturn. The wages of job stayers in CH-firms were less constrained by the minimum wage than hiring wages, since stayers are generally paid more than new hires (see Figure A.1D).

To answer the question of how hiring wages might have responded to the Great Recession in the absence of a binding minimum wage, we use the kernel re-weighting method of DiNardo et al. (1996). A description of this method is provided in Appendix A.6. A partial equilibrium assumption underlies this method: the number and composition of entry-level jobs is not affected by the NMW. This assumption is unlikely to hold in reality. Nevertheless, this method allows us to assess the impact of the NMW in a simple and transparent way.

Here we briefly explain the intuition. For each year following 2004, we replace the density of job-level real hiring wages which was at or below the real value of the NMW in that year, with the corresponding section of the 2004 density, adjusted for differences in observable job characteristics. Then, we re-scale this counterfactual density so that the two sections integrate to one. We select 2004 as the base year because this was the last year when the real value of the NMW was below its lowest level in 2014 (see Figure 1.3). For this estimation we use the plug-in method of Sheather and Jones (1991) to select the optimal bandwidth, which ranges from 0.01 to 0.04 for our sample.

The most important parameter in this kernel re-weighting exercise is the assumed size of the spillover effect of the minimum wage, i.e. the highest value of the real

[10]The adult rate age limit was decreased from 22 to 21 in 2010.

FIGURE 1.3: Real hourly wages of new hires and NMW adult rate, ages 22-64

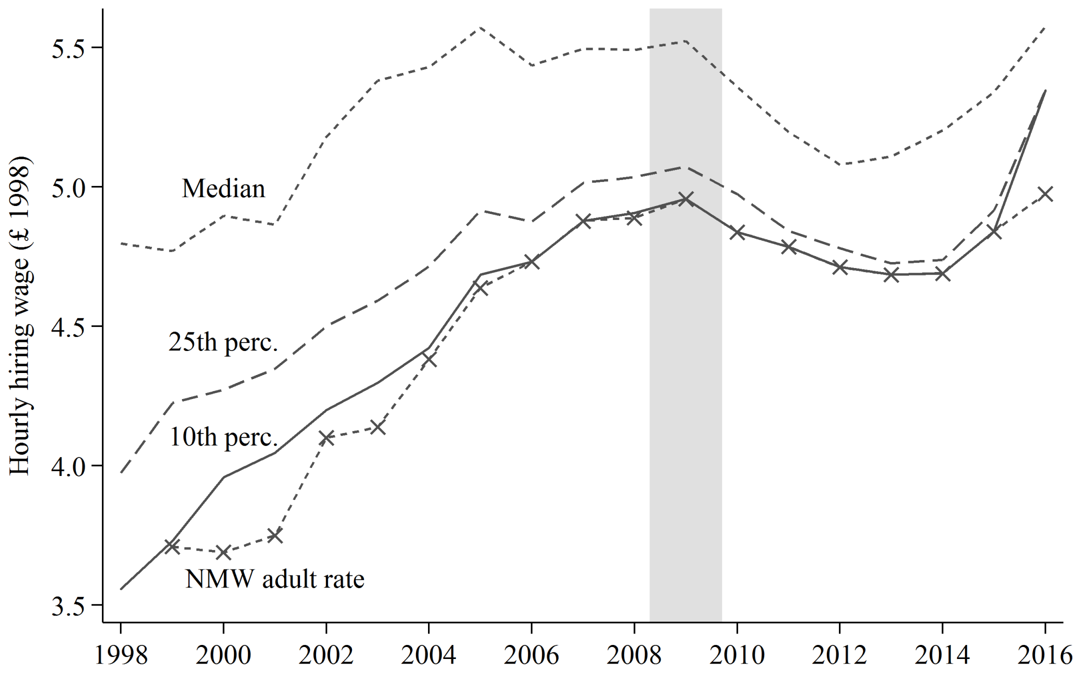

Notes.- National Minimum Wage adult rate and 10th, 25th, and median percentile of job-level hourly hiring wages, ages 22-64. All monetary values are deflated to 1998 values using the CPI. Shaded area marks official UK recession.

hiring wage density which is affected by the NMW. The more spillover we assume in a period, the more of this period's density - the section below the real value of the minimum wage plus any spillover - is replaced with the corresponding section of the 2004 density.

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

To compute hiring wages at the job-level from counterfactual densities, we assume that the rank of a job in the distribution of hiring wages is preserved under different values of the NMW. Then we re-estimate regressions (1.1) and (1.2), using each of the counterfactual real hiring wage samples estimated with different spillover parameters. Figure 1.4 displays the point estimates of the counterfactual semi-elasticity of real hiring wages with respect to the unemployment rate across a range of assumed parameters of the spillover.

FIGURE 1.4: Counterfactual estimates of the semi-elasticity of real hiring wages, 1998-2016: varying the assumed spillover effect of the NMW

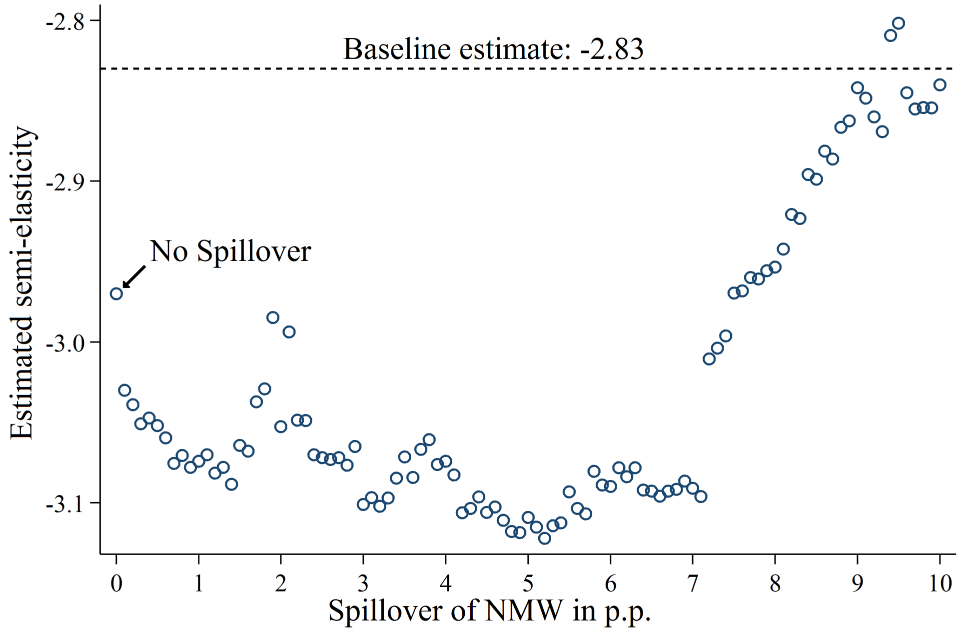

Notes.- each circle represents an estimate of the semi-elasticity of real hiring wages with respect to the unemployment rate. Standard errors lie outside of the figure. The horizontal axis shows the assumed spillover effect in p.p.. Dashed line shows the baseline estimate of semi-elasticity. We use a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

Assuming that there is no spillover effect, the left-most circle shows that the responsiveness of real hiring wages to the unemployment rate increases from -2.83 to -2.97 per cent. The standard errors are comparable to the baseline value (0.9) and lie outside the range of this figure. The semi-elasticity falls below -3.1 when the spillover effect increases to five p.p. If the spillover increases above seven p.p., then the responsiveness of real hiring wages to changes in the unemployment rate begins to decrease towards the baseline estimate: the shape of the counterfactual density increasingly resembles the shape of the density observed in 2004 when we assume larger spillover effects, and hence the variation of hiring wages over time in entry-level jobs declines. These results suggest that the NMW constrained firms in how far they could reduce wages of new hires during the Great Recession.

1.5 How did wages and hours evolve after hiring?

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

As an initial assessment of the importance of cohort effects, Figure 1.5 plots the real hourly wages and hours worked averaged over employees instead of jobs, for each cohort of entry-level new hires, conditional on these employees staying in their respective jobs. The average hiring wages and hours in each year are shown as solid lines. Panel A suggests that wages exhibit cohort effects: the real wages of hiring cohorts from 1998 to 2005 mostly seem to have parallel trends in the first three years on the job, similar to the findings of Baker et al. (1994) for one US firm. But, unlike these authors, we see that the cohort-specific paths of wages respond to the business cycle, as shown by a decline in wage growth during the years of the Great Recession. Cohorts hired during this time seem to be locked into low wage growth trajectories. For example, the mean wages of the 2013 cohort in 2015 were still below the mean wages of the 2014 cohort in 2015. Panel B of Figure 1.5 similarly suggests that the path of hours worked depends on cohort effects, though less strikingly so than for wages, as growth trends remained mostly parallel throughout the period. In other words, differences in cohort hiring wages and hours over the business cycle seem to persist, and may even reinforce the initial decline in labour costs.

Comparing sample averages over time is likely to be subject to a composition bias, since relatively low-wage employees are given less weight during downturns than in normal times (Solon et al., 1994). Therefore we estimate how the wages and hours of new hires in entry-level jobs evolved over three years of subsequent tenure, including match-fixed effects to control for the changing composition of matches over the business cycle. We include only consecutive observations of a worker in some job, such that a worker with three years of tenure must be observed the previous two years.

FIGURE 1.5: Paths of real wages and hours for cohorts of new hires

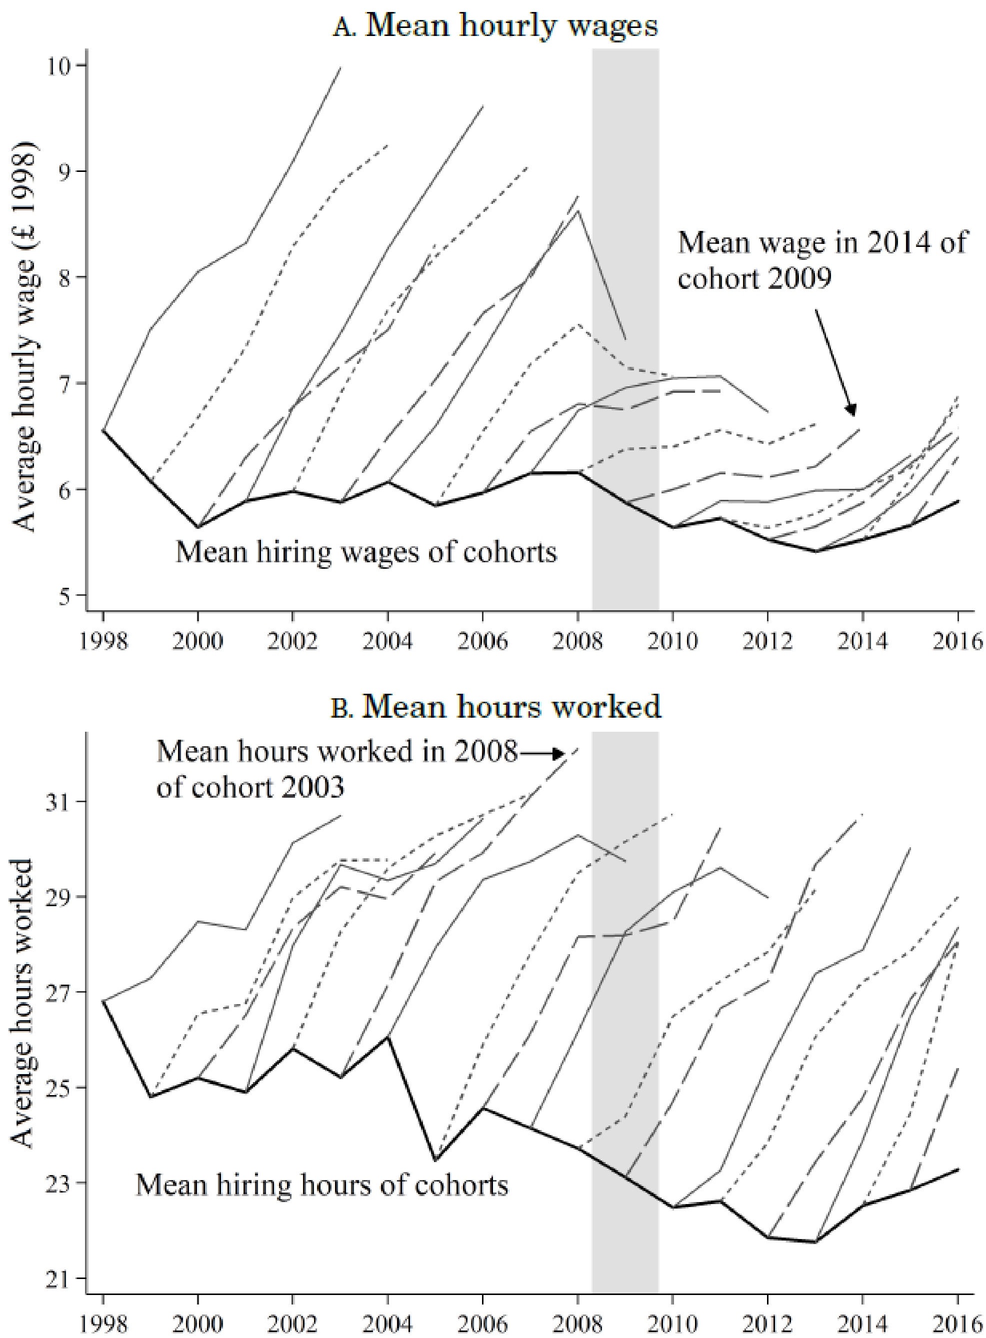

Notes.- the solid lines give the average real hiring wage and weekly hours worked for each cohort of new hires in our sample of entry-level jobs (i.e. column (1), Table 1.1). Each line branching off from the solid line shows the paths of wages or hours of these hiring cohorts over time, as their tenure in the job increases. When employees leave their hiring jobs they also exit the samples of their respective cohorts.

The sample of workers for each hiring year is unbalanced, since workers exit from entry-level jobs: either they switch jobs within the same firm or across firms, or they exit into non-employment. Using least squares we estimate

$$w_{m\tau} = \theta_m + \psi_{S(m)\tau} + x'_{m\tau}\phi + \eta_{m\tau} \ , \tag{1.4}$$

where the dependent variable is the log real wage of some match m between a worker (i) and job (j) with tenure τ, where $S(m)$ is a function indicating in which year s a match was formed. θ_m is a match-fixed effect and $\psi_{s\tau}$ are cohort-tenure-fixed effects for any matches beginning in year $s = 2001, ..., 2013$ with years of tenure $\tau \in [0,3]$. The sample size per cohort is initially over a thousand employees with tenure greater than a year, and then declines to around four hundred employees per cohort with tenure over three years. The vector $x_{m\tau}$ contains time-varying quadratic controls for the size of the firm, and $\eta_{m\tau}$ is the error term. We estimate (1.4) by excluding the effects ψ_{s0}, so the estimated values of $\hat{\psi}_{s\tau}$ for $\tau > 0$ are interpreted as log changes relative to the hiring wages in entry-level jobs. Although there is certainly endogenous selection of the employees who stay in these jobs for up to three years, the match-fixed effect should partially address this concern. We similarly estimate (1.4) with log basic weekly hours worked as the dependent variable.

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

We plot the estimated cohort-tenure-fixed effects in Figure 1.6 for selected hiring years, and the underlying estimates are displayed in Appendix Table A.9 for all years with confidence intervals. The last cohort of hires unaffected by the Great Recession within three completed years of tenure was 2005. Panel A shows a clear U-shaped response of real wage growth at all levels of tenure over the Great Recession. Over the first year on the job, the wages of workers hired in 2009 were stagnant, while for those hired before the recession they grew on average by two percent, and for those in the last cohort by three per cent. Similarly, there was no real wage growth over the subsequent three years for 2009 hires, compared with over eight per cent for 2005 hires.

The cyclical differences in hours growth in these jobs are less pronounced. Initially the increases over three years were smaller in 2008 than pre-recession, with negative growth over the first year of tenure. However the average increases following the recession were greater. This latter period coincides with the persistent rise and peak

FIGURE 1.6: Estimated composition-adjusted and cohort-specific log changes $\hat{\psi}$ in real wages and hours relative to hiring levels: workers who stay in entry-level jobs

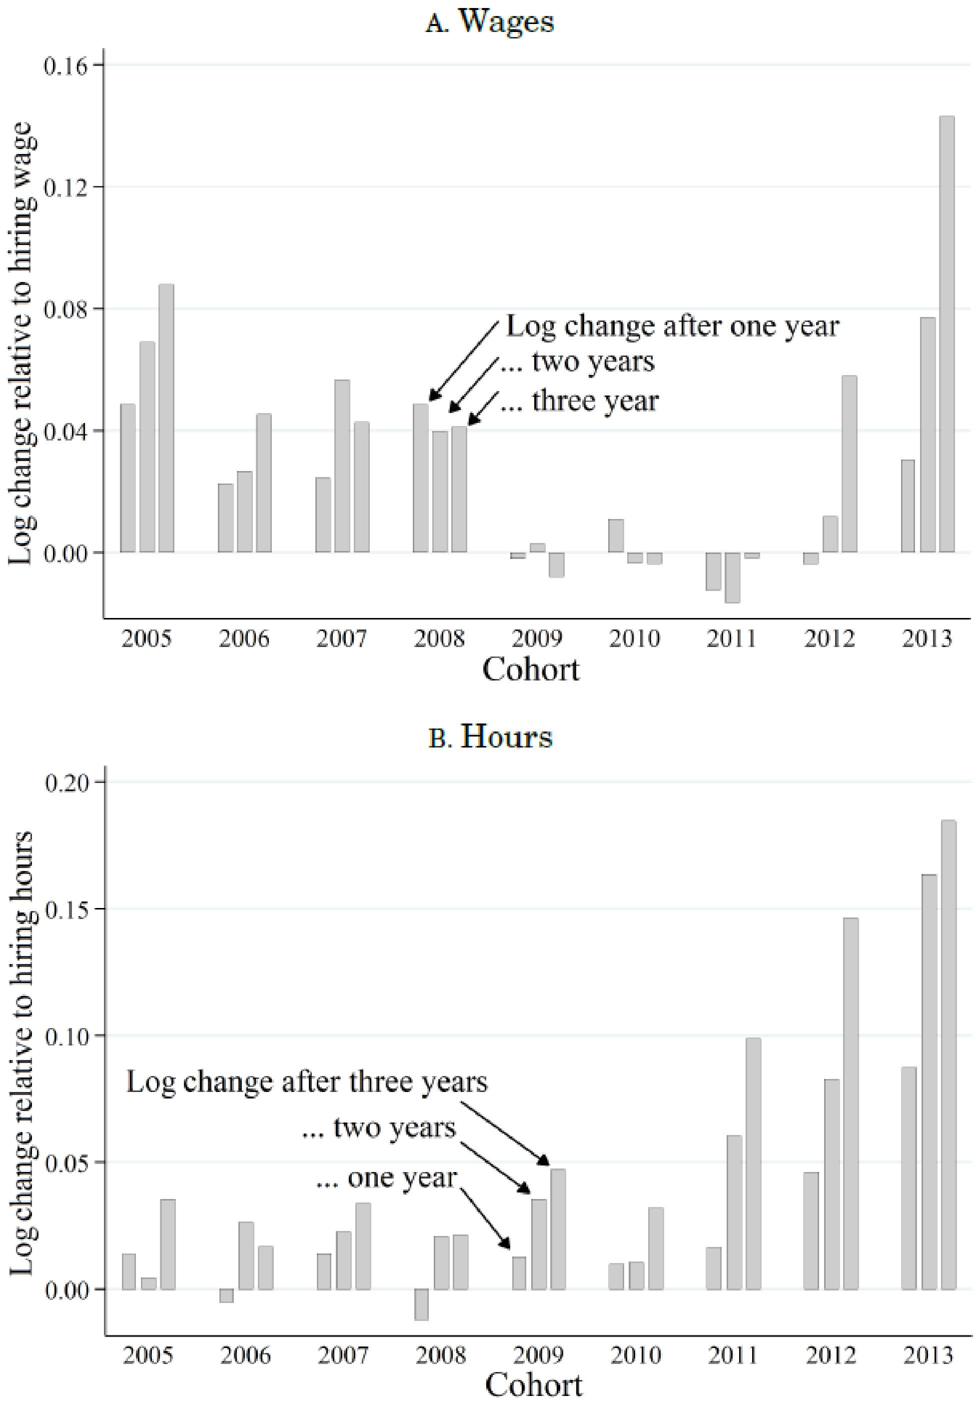

Notes.- cohort average change in log wages and log hours with tenure, relative to respective hiring values, in entry-level jobs. Composition adjusted by controlling for match-fixed effects. See Appendix Table A.9 for standard errors and results for all other hiring years in 2001-13.

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

The findings in this section suggest that firms were not only able to significantly reduce the real wages and hours of new hires in response to the Great Recession, but also depress wage growth with subsequent tenure. However there are at least two reasons this evidence is only suggestive. First, it only applies to workers who stay in the exact same job in the firm, whereas in reality, expected employee progression or reallocation to other jobs within the firm also affects the ex ante present value of a match and the hiring decision. Second, the regression in (1.4) is subject to the same measurement criticism which the majority of this paper shows is important: it does not control for the endogenous selection and weighting of matches over time, which we are unable to adequately address due to a small number of degrees of freedom at the job level here.

1.6 Discussion of findings and possible explanations

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

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

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

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

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

Montgomery (1988) discusses the factors which determine firms' demand for part-time employees. If there are fixed costs of hiring and training new employees, then these costs are unlikely to vary between part- and full-time hires in the same job: the ratio of hours to fixed costs will often be lower for part-time hires. Firms require compensation for this lower return from part-time hiring, that is, the hourly wage per worker has to be lower. This firm-side compensating differential should be stronger for higher-skilled jobs, where hiring and training costs are typically greater. Montgomery (1988) provides evidence for these features of wage-setting and hiring behaviour in the presence of fixed costs among US establishments. Moreover, if firms have to pay all workers in some job the same hourly rate, then firms are more likely to employ full-time employees when there are fixed hiring costs. However, fringe benefits (pension contribution, health care) function as quasi-fixed costs which might only be offered to full-time employees, and thus shift the demand from full- to part-time

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

workers. To the extent that these fixed costs depend on the level of productivity, it is possible that they decline during recessions, and thus make part-time hiring more likely. The cyclical properties of fixed hiring costs in the UK is an interesting empirical question for future research.

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

1.7 Conclusion

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

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

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

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

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

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

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

Chapter 2

Unemployment and econometric learning

Note: This chapter will appear as an article in Research in Economics; DOI: 10.1016/j.rie.2017.10.005. An older version has also been published as Edinburgh School of Economics Discussion Paper Series Nr. 267. This article was co-authored with Carl Singleton, who is a Post-doctoral Fellow at the University of Edinburgh, School of Economics; e-mail: carl.singleton@ed.ac.uk. Carl has agreed that this essay represents in the majority my work, and that it can appear within this thesis.

2.1 Introduction

The Mortensen and Pissarides (1994) model of search and matching frictions has become the foundation for studying the cyclical behaviour of labour markets (see Rogerson and Shimer, 2011 for a survey). The existence of search frictions in labour markets is usually motivated by decentralisation, due to geography and other differences between firms, and because each worker has characteristics which make them more or less suitable for the available jobs. While they search and decide whether or not to accept a match at some agreed upon wage, unemployed workers and firms must form expectations of future variables relevant to their choices, including aggregate conditions. The decentralised nature of labour markets makes it a priori *not* obvious that workers and firms are able to correctly forecast these variables at all times. Nonetheless the rational expectations (RE) assumption is usually made in search and matching models. This is likely to pose overly strong requirements on the cognitive abilities of economic agents, and is unrealistic in the presence of potentially frequent structural or policy shocks. Even small departures from RE might alter the qualitative or quantitative predictions of these models.

Here we analyse the equilibrium properties and dynamics of the textbook real business cycle (RBC) Mortensen-Pissarides model (see Hagedorn and Manovskii, 2008 for the standard discrete-time treatment), while representing agents as ‘good econometricians’, who form forecasts according to their estimates of some structural model parameters (Evans and Honkapohja, 2001).[1] We assume that agents employ a recursive least squares (RLS) algorithm to update their parameter estimates when new data become available. Econometric (or adaptive) learning can provide a behavioural foundation for RE if the rational expectations equilibrium (REE) is shown to be learnable or E-stable (Evans and Honkapohja, 2001); i.e. small deviations from this equilibrium are reversed over time, with asymptotic convergence. We show that the model’s unique REE is E-stable. No parameter restrictions are required to ensure this is the case, beyond those which make the model well-formulated, when agents use a minimum-state-variable rule to form and update their forecasts of so-called labour market tightness. Furthermore, we confirm that this equilibrium is globally stable and satisfies the properties of Strong E-stability (i.e. being robust to over-parametrisation of the econometric relationship by the agents). And so from this perspective, the assumption of RE when studying or applying this model would seem to be reasonable.

This article contributes to a significant literature on the more realistic representation of agents as behaving like econometricians in macroeconomic models. Mankiw et al. (2004) offer empirical evidence against RE. Their analysis of surveys of professional forecasters and households finds significant autocorrelation in forecast errors, which is compatible with econometric learning, but not with RE. Milani (2007, 2011) also argues for the presence of adaptive learning in the New Keynesian model. He shows that learning by agents is capable of replacing the other ‘mechanical’ sources of persistence in these models, such as habits, whilst at the same time increasing the fit to the data as compared to assuming RE. Pfajfar and Santoro (2010) examine a survey of households’ inflation forecasts over several decades and conclude that the hypothesis of RE can be rejected, and that there is evidence in support of adaptive learning dynamics. This view is further supported by Berardi and Galimberti (2012), who examine post-WWII US inflation and output growth. Comparing the performance of different adaptive learning algorithms in matching survey forecasts, their results suggest that economic agents form these according to RLS.

[1] Strictly, the characterisation of agents who behave in this way as ‘good econometricians’ is used by Branch and McGough (2016). One should think of these as being agents who make conditional forecasts, which are pertinent to their decisions and based on a simple model such as linear regression, and who update this model based on forecast errors.

The formulation of learning we use here is such that agents need only make one-step-ahead forecasts of the labour market's condition. Ours is not the only recent study to apply the principles of econometric learning to this class of model. From a similar set-up, Di Pace et al. (2016) consider an approach where agents must make infinite horizon forecasts about the future paths of wages, unemployment and profits in order to make choices today. This latter type of learning fits into the anticipated utility approach (Kreps, 1998), and has notably been applied to the RBC model by Eusepi and Preston (2011). Di Pace et al. focus on results with constant gain learning, for which there are no equivalent analytical results to the E-stability conditions we consider. The authors use the model to address the 'unemployment volatility puzzle': the inability of the Mortensen-Pissarides model to generate a realistically large amplification of unemployment for a given change in wages or productivity (Shimer, 2005). Under infinite horizon learning, they not only match US professional forecast errors, but also find a greater cyclical unemployment rate response relative to the baseline model. This is driven by persistence or inertia in agents' expectations of the future path of wages, which implies that firms are over-optimistic about future profits, post more vacancies, and thus unemployment is more volatile relative to the REE baseline case. They also find some, but significantly less, propagation of the unemployment response when the model is reformulated in a one-step-ahead forecast guise. Kurozumi and Van Zandweghe (2012) also apply econometric learning to an extended model of the business cycle, which includes sticky prices and monetary policy, as well as labour market search frictions. They analyse determinacy and E-stability conditions, finding that these depend on model parameters. However, both this set-up and that of Di Pace et al. (2016) differ form our own in so far as they move beyond the textbook Mortensen-Pissarides model, and in both cases the agents must forecast several aggregate variables, which do not appear in the reduced form of the REE characterisation.

We also present illustrative simulations and analyse the dynamics of the unemployment model, and show that the REE could be a poor approximation to an economy in which agents are econometricians. Convergence to the REE is very slow, even when agents have a short memory or give greater weight to more recent data. If agents must learn the REE, then aggregate variables may be persistently some distance away from their REE equivalents. We also demonstrate how structural shocks generate a more gradual cyclical adjustment of wages after the introduction of learning, and thus predicted unemployment volatility when such shocks occur frequently would be reduced relative to the REE. Therefore we find some different results to Di Pace et al. (2016). The same amplification mechanism described therein is not present here. By keeping closer to the spirit of the most standard RBC variant of the model, in which the only relevant choice is the number of vacancies that firms post,[2] agents need only estimate a relationship between labour market tightness and productivity to form expectations and close the model, and so wage determination is absorbed. Econometric learning in our set-up then generates inertia in expectations of tightness (and wages) following shocks.[3]

The REE of the model describes a choice of labour market tightness which is independent of the state of unemployment. Therefore, agents' learning of the minimum-state-variable solution implicitly assumes their complete understanding of the economy's dynamics of unemployment and vacancy creation. We consider an alternative decision rule, which relaxes this latter implied assumption. We consider whether or not agents can learn how many vacancies to create in response to an expanded state of the world, which includes the state of unemployment. In other words we also ask if agents can learn the Beveridge curve. This alternative is not E-stable for the complete range of possible model parameters. And where the economy does converge to the REE, it does so more slowly than under the minimum-state-variable representation. In this respect, the approximation of the REE for this model would be even further weakened.

2.2 The search and matching model of unemployment

We outline a discrete-time search and matching model of the labour market, following Merz (1995) and Andolfatto (1996) by assuming that members of a representative household perfectly insure each other against income fluctuations. We derive a difference equation for so-called labour market tightness which summarises its equilibrium, thus analogous to the treatment in Hagedorn and Manovskii (2008).[4] This can be regarded as a textbook model, which has been applied, critiqued and extended exhaustively in the literature, not least in attempts to solve the unemployment volatility puzzle.

[2]With linear production technology this is also equivalent to the decision of vacancy creation or destruction when firms consist of a single worker. Given worker homogeneity, we will rule out states of the world whereby workers would choose not to work.

[3]Although we do not expand on this point later, it is straightforward to see that the minimum-state-variable solution of the model we apply, which agents learn and use to form expectations, could be re-written in terms of wages and productivity by substituting for the standard 'wage curve' derived from Nash bargaining, (2.19).

[4]Our set-up of the model only differs in so far as we describe a representative household and firm, rather than a continuum of the latter. The characterisation of the equilibrium is approximately identical, though we hope our exposition is more familiar as an extended RBC model with aggregate uncertainty.

2.2.1 The labour market

There is a continuum of identical, risk-neutral workers with total measure one, and an infinite horizon. The matching function $M(u_t, v_t)$ provides the number of successful matches in period t. It is increasing and concave in both of its arguments, u_t and v_t, which represent the share of the total workforce currently unemployed and the level of vacancies relative to the size of the workforce respectively. Matches and separations occur after agents in the economy have made decisions, i.e. at the end of each period. A Cobb-Douglas, constant returns to scale matching function is chosen, due to its simplicity, well-known features and being most common in the literature:

$$M(u_t, v_t) = \mu u_t^{\alpha} v_t^{1-\alpha}, \qquad \mu > 0, \qquad \alpha \in (0,1), \tag{2.1}$$

where μ gives a measure of matching efficiency and α the elasticity of the number of matches with respect to unemployment. We define the level of labour market tightness as

$$\theta_t = \frac{v_t}{u_t}. \tag{2.2}$$

Unemployed workers and vacancies are matched randomly, and so the probability of a firm filling an open vacancy each period is

$$q(\theta_t) = \frac{M(u_t, v_t)}{v_t} = \mu \theta_t^{-\alpha}. \tag{2.3}$$

The corresponding probability that an unemployed worker gets matched to an open vacancy is

$$\theta_t q(\theta_t) = \frac{M(u_t, v_t)}{u_t} = \mu \theta_t^{1-\alpha}. \tag{2.4}$$

Matches are destroyed with probability $\lambda \in [0,1]$. Unemployment changes between periods due to inflows, caused by exogenous separations, and outflows from new matches. The resulting law of motion for the share of workers unemployed at the beginning of period $t+1$ is

$$u_{t+1} = u_t + (1-u_t)\lambda - \theta_t q(\theta_t) u_t, \tag{2.5}$$

whereby the final two terms measure these inflows and outflows respectively. The steady-state level of unemployment, where these flows are equal, for given θ_t is

$$u_t^* = \frac{\lambda}{\lambda + \theta_t q(\theta_t)}. \tag{2.6}$$

2.2.2 The household

For expositional simplicity we consider an economy comprising a single infinitely-lived representative household of size one, in which all workers are identical and risk neutral. There is perfect consumption insurance across its members. In period t, n_t of the household's members are employed and $(1-n_t)$ are waiting for a match. In other words, labour supply is inelastic, which can be ensured by assuming that the equilibrium wage is always strictly greater than the per-period utility value from non-employment b; formally we assume $w_t > b \geq 0$. The household discounts each additional period's utility by a factor $\delta \in (0,1)$. In period t the household has risk-neutral preferences over income a_t and has expected lifetime utility

$$E_t^*\left[\sum_{s=0}^{\infty} \delta^{t+s} a_{t+s}\right], \tag{2.7}$$

where E_t^* denotes expectations (not necessarily rational). Income is given by

$$a_t = w_t n_t + b(1-n_t) + D_t, \tag{2.8}$$

with $w_t n_t$ and $b(1-n_t)$ denoting total income from labour and non-employment respectively, and D_t are dividends from a representative firm, which is owned by the household. This household can also be represented by the Bellman equation

$$W(n_t) = w_t n_t + D_t + b(1-n_t) + \delta E_t^*\left[W(n_{t+1})\right], \tag{2.9}$$

where $W(n_t)$ represents the household's current value function, with their state of the world given by the employment level $n_t = 1-u_t$. The household takes as given wages w_t, dividends from the representative firm D_t and labour market tightness θ_t. The household's expected continuation value is $E_t^*\left[W(n_{t+1})\right]$. The law of motion for employment follows directly from (2.5), and is given by

$$n_{t+1} = (1-\lambda)n_t + \theta_t q(\theta_t)(1-n_t). \tag{2.10}$$

Applying the envelope theorem to (2.9) and using the law of motion (2.10), the marginal value of household employment is given by

$$\frac{\partial W(n_t)}{\partial n_t} = w_t - b + \delta(1-\lambda-\theta_t q(\theta_t))\frac{\partial E_t^*\left[W(n_{t+1})\right]}{\partial n_{t+1}}, \tag{2.11}$$

i.e. the net utility value from wages exceeding non-employment income plus the discounted expected continuation value from additional employment.

2.2.3 The firm

The production side of the economy consists of a representative firm. This firm employs workers and produces output $y_t n_t$, where y_t is the marginal product of labour. We assume that the process of worker productivity is a stationary AR(1) process in logs:

$$\log(y_t) = \rho \log(y_{t-1}) + \varepsilon_t, \qquad \rho \in [0,1), \tag{2.12}$$

with some initial condition y_0 and ε_t being drawn as an *iid* zero-mean shock. Because labour is the only input into production, the firm maximises profits by choosing the level of employment, subject to the law of motion for the labour market. However, due to the exogeneity of separations, the optimal choice for the level of employment and the optimal quantity of vacancies to open coincide. For each vacancy held open the firm has to pay per-period unit cost $c > 0$.

The objective of the firm is to maximise the expected value of current and future profits, given by

$$E_t^* \left[\sum_{s=0}^{\infty} \delta^{t+s} \left(y_{t+s} n_{t+s} - w_{t+s} n_{t+s} - c v_{t+s} \right) \right]. \tag{2.13}$$

Because the firm is owned by the household, the firm discounts future profits using the same discount factor as the household. The firm takes wages w_t and labour market tightness θ_t as given. Employment at the firm follows the law of motion

$$n_{t+1} = (1-\lambda) n_t + q(\theta_t) v_t, \tag{2.14}$$

which shows that the more vacancies the firm creates, the higher aggregate employment will be in the next period. Maximising (2.13), subject to (2.14), can be represented as the Bellman equation

$$\Pi(v_t; n_t, y_t) = \max_{v_t \geq 0} \left(y_t n_t - w_t n_t - c v_t \right) + \delta E_t^* \left[\Pi(v_{t+1}; n_{t+1}, y_{t+1}) \right], \tag{2.15}$$

where $\Pi(v_t; n_t, y_t)$ represents the firm's current value function. Profit maximisation in this case implies that the representative firm will open or close vacancies until the marginal cost and benefit of doing so are equal:

$$\frac{\partial E_t^* \left[\Pi(v_{t+1}; n_{t+1}, y_{t+1}) \right]}{\partial n_{t+1}} = \frac{c}{\delta q(\theta_t)}. \tag{2.16}$$

Applying the envelope theorem and using the above first order condition gives the surplus to the firm from employing an additional worker,

$$\frac{\partial \Pi(v_t; n_t, y_t)}{\partial n_t} = y_t - w_t + \frac{(1-\lambda)c}{q(\theta_t)}, \tag{2.17}$$

i.e. the net profit from employing an additional worker plus the discounted expected continuation value, taking matching frictions into account. The optimal choices of the firm (2.16) and (2.17) imply that labour market tightness evolves according to the non-linear difference equation

$$\frac{c}{\delta q(\theta_t)} = E_t^* \left[y_{t+1} - w_{t+1} + \frac{(1-\lambda)c}{q(\theta_{t+1})} \right]. \tag{2.18}$$

In other words, the representative firm must form expectations about the right hand side of (2.18) to optimally choose the number of vacancies to open in the current period. In particular, the firm forecasts labour productivity y_{t+1}, the real wage which will be realised next period w_{t+1}, and the value of labour market tightness in the next period θ_{t+1}. Note that this problem does not depend on the type of expectations formation; we have not specified how forecasts of the right hand side of (2.18) are formed.

2.2.4 Wage determination

Wages are determined by generalised Nash bargaining between the firm and workers over the additional surpluses (2.11) and (2.17), with worker bargaining power $\beta \in [0,1]$:[5]

$$w_t = \arg\max \left(\frac{\partial W(n_t)}{\partial n_t} \right)^{\beta} \left(\frac{\partial \Pi(v_t; n_t, y_t)}{\partial n_t} \right)^{1-\beta}.$$

Combining the surplus sharing rules which form the solution of this problem, iterating forwards, and using (2.11), (2.16) and (2.17) gives what is referred to in the textbook model as the ‘wage curve’:

$$w_t = (1-\beta)b + \beta(y_t + c\theta_t). \tag{2.19}$$

To ensure employment is always preferred and a wage successfully negotiated we also restrict $y_t > b$. Again, we do not have to specify rational expectations to obtain the wage curve.

[5]Note, it is crucial here that both workers and firms are assumed to form expectations in the same way, using the same rule, as otherwise the Nash bargaining solution would be significantly complicated. In the sense of the model here, since workers own the firm, this is not an unreasonable assumption.

2.2.5 The rational expectations equilibrium

Our search and matching framework consists of the goods and the labour market. Since our focus is to study the labour market under econometric learning, we abstract from the goods market. This can be justified by Walras' Law, which states that equilibrium in the labour market implies that the goods market clears.

The representative household supplies labour inelastically, and thus the REE of the model can be summarised and determined uniquely by the value of labour market tightness at which point the representative firm is indifferent between opening an additional vacancy or not. In other words, the firm has to form expectations about the future state of labour market tightness as this affects the current discounted value of a match. The non-linear difference equation determining this value of θ_t, substituting the outcome of the wage bargaining (2.19) into (2.18), is given by

$$\frac{c}{\delta q(\theta_t)} = E_t^* \left[(1-\beta)(y_{t+1} - b) + \frac{(1-\lambda)c}{q(\theta_{t+1})} - \theta_{t+1}\beta c \right]. \tag{2.20}$$

To provide intuition for this expression we stress the similarities to (2.18). Aggregate labour market tightness θ_t will adjust immediately to deviations from this equality via the firm instantaneously opening or closing vacancies. Thus today's labour market tightness is determined by expectations of the value of a filled vacancy in the next period. In equilibrium it must also be that $D_t = y_t n_t - w_t n_t - c v_t$.

Given the process for productivity (2.12), the equilibrium θ_t is the solution of the non-linear difference equation (2.20). With this and initial condition u_1, the remainder of the interesting endogenous variables in the equilibrium, $\{a_t, w_t, v_t, u_{t+1}\}_{t>0}$, can be obtained using (2.2), (2.5), (2.8) and (2.19). In the next section we linearise around steady-state values to obtain an analytical solution to (2.20) and discuss the rational expectations equilibrium.

2.3 Linearisation and the rational expectation equilibrium

To solve the system consisting of (2.12) and (2.20), we linearise around deterministic steady-state values $\bar{\theta}$ and $\bar{y} = 1$:[6]

$$\theta_t = \psi_0 + \psi_1 E_t^* y_{t+1} + \psi_2 E_t^* \theta_{t+1}, \tag{2.21}$$

$$y_t = (1-\rho) + \rho y_{t-1} + \varepsilon_t, \tag{2.22}$$

where the coefficients are functions of the model's parameters and steady state values

$$\begin{aligned}
\psi_0 &= [1-\psi_2]\,\bar{\theta} - \psi_1 \bar{y},\\
\psi_1 &= (1-\beta)\delta\bar{\theta}q(\bar{\theta})(c\alpha)^{-1},\\
\psi_2 &= \delta\left[(1-\lambda) - \beta\bar{\theta}q(\bar{\theta})\alpha^{-1}\right].
\end{aligned}$$

We now assume that expectations are rational, that is, the firm and the household take all available information into account and forecast θ without systematic errors. We denote the rational expectations operator by E_t. Linear RE models where agents form expectations regarding an endogenous variable can have multiple equilibria (or *bubble* solutions, not related to economic fundamentals).[7] If there are multiple *stable* REEs then a model is said to be indeterminate. In B.1.2 we show that a unique stable equilibrium exists so long as $|\psi_2| < 1$. Intuitively, this condition requires that the shocks ε are transitory and θ returns to its steady state value, analogous to a stationary first-order autoregressive process. The solution in this case to the system (2.21) and (2.22) is obtained by using the method of undetermined coefficients. After substituting $E_t y_{t+1} = (1-\rho) + \rho y_t$ in (2.21), the solution can be guessed to have the form

$$\theta_t = A + B y_{t-1} + C\varepsilon_t, \tag{2.23}$$

since the only predetermined variable in the above system is y_{t-1}. The parameters A, B, and C of the reduced form are functions of the parameters in (2.21) and (2.22).[8]

[6] See B.1.1 for derivation and all subsequently defined parameters, such as ψ's, expressed in terms of the model parameters and steady-state values.

[7] The literature on bubbles and the related concept of indeterminacy is reviewed in Benhabib and Farmer (1999). The classic reference on bubble solutions is Blanchard and Watson (1983); see also Bullard and Mitra (2002) for an analysis of indeterminacy in a New Keynesian framework.

[8] See B.1.3 for a description of how (2.23) can be obtained from (2.21) and (2.22), and also A, B and C in terms of the model parameters.

Using equation (2.23), RE about next period's labour market tightness are given by

$$E_t \theta_{t+1} = A + B y_t, \tag{2.24}$$

since $E_t \varepsilon_{t+1} = 0$. The parameters used thereby to estimate θ_{t+1} are true values, that is the firm and the household know the true underlying functional forms and their associated parameter values. When ε_t is realised in period t, it becomes part of the information set of firms and households and the resulting forecast $E_t \theta_{t+1}$ leads to an immediate adjustment of vacancies, such that the difference between $E_t \theta_{t+1}$ and θ_{t+1} is only from the next period's shock $C\varepsilon_{t+1}$. In the absence of new shocks the firm and the household forecast the response of market tightness to productivity correctly using (2.24).

2.4 Adding econometric learning to the model

We now depart from RE and apply the concept of econometric learning to this model. Unlike the application of Di Pace et al. (2016), who suggest that agents might need to form infinite horizon forecasts of multiple variables, such as wages or firm profits, we depart from the REE summarised above by the single choice variable θ_t, which just requires a one-step-ahead forecast of θ_{t+1}. Kurozumi and Van Zandweghe (2012) have also considered the role of learning when agents need to make one-step-ahead forecasts in the presence of labour search frictions. However, their model also includes sticky prices and monetary policy, involves forecasting over several aggregate variables, and the determinacy of the REE is not always certain. In what follows the expectational stability results become more clear-cut.

2.4.1 E-Stability of the MSV solution

We relax the assumption of RE by modelling agents as econometricians attempting to estimate the parameters A, B, and C, which underpin the true motion of the economy under uncertainty. Agents are endowed with a perceived law of motion (PLM) in the economy of the MSV form (2.23), because we derived its functional form without having to impose rational expectations.[9] In other words, agents know the structure of the economy as expressed in the system (2.21) and (2.22), but the parameter values are unknown to them. They make corresponding estimates of the true coefficients in period t, given by $\hat{A}_t$, $\hat{B}_t$ and $\hat{C}_t$, and update these each period when new data becomes

[9] If agents are not able to learn the simplest representation (as few state variables as possible), they cannot be expected to learn equilibria containing more state variables and to coordinate behaviour towards them.

available. Therefore, the household and firm forecast as under RE in (2.24), but instead of the true parameter values they use estimates $\hat{A}_t$ and $\hat{B}_t$ to forecast labour market tightness

$$E_t^* \theta_{t+1} = \hat{A}_t + \hat{B}_t y_t. \tag{2.25}$$

The main difference is that forecasts with rational expectations coincide with the true realisation of next period's θ on average, whereas this is not the case for econometric learners. They possess less information than rational agents, since they do not know the parameters of the model. We assume econometric learners perform the task of estimating parameters using recursive least squares (RLS). This is the most widely used estimation technique in the learning literature and Berardi and Galimberti (2012) provide evidence that this estimator matches surveys of forecasts of US time series closely.[10] Let the vector of parameter estimates be denoted by $x'_{t-1} = (\hat{A}_t, \hat{B}_t)$, then the general recursive updating algorithm can be represented by

$$x_t = x_{t-1} + g_t Q(\theta_t - \hat{A}_t - \hat{B}_t y_{t-1}), \tag{2.26}$$

which shows that agents update their previous parameter estimates x_{t-1} by a function of the observed forecasting error. The function Q and the so-called "gain parameter" g_t are further described below. There are potential problems of simultaneity in forward looking models. Therefore, it is assumed that although agents forecast θ_{t+1} using y_t, the variable y_t is *not* in the information set for the estimation of $\hat{A}_t$ and $\hat{B}_t$. As proved by Marcet and Sargent (1989), this does not alter the asymptotic stability results obtained in the following, as compared to an algorithm allowing for simultaneity, so long as agents are assumed to ignore outliers, defined as being observations outside of some predetermined range. This timing assumption is usually thought of as realistic, since robust macroeconomic data is only available to decision makers with a substantial lag.[11]

Since the current value of labour market tightness depends on the prediction of next period's value, agents estimates have the potential to affect the path of labour market tightness. To see this, we substitute the stochastic process of labour productivity and the econometric forecast (2.25) into (2.21), which gives the actual law of motion

[10] The presented algorithm is comparable to a restricted form of the Kalman filter. For further discussion see Berardi and Galimberti (2013).

[11] The assumption also plausibly implies that under subjective expectations agents would only ever enter the wage bargaining process with pre-determined valuations. Otherwise, there would be simultaneity between the bargaining result and subsequent expectations formation.

(ALM) for labour market tightness:

$$\begin{aligned}\theta_t =& \psi_0 + \psi_1(1-\rho)(1+\rho) + \psi_2\hat{A}_t + \psi_2\hat{B}_t(1-\rho) \\ &+ (\psi_1\rho + \psi_2\hat{B}_t)\rho y_{t-1} \\ &+ (\psi_1\rho + \psi_2\hat{B}_t)\varepsilon_t. \end{aligned} \tag{2.27}$$

This defines the following T-mapping from the PLM, $\theta_t = A + By_{t-1} + C\varepsilon_t$, to the ALM:

$$\begin{aligned} T(\hat{A}_t) &= \psi_0 + \psi_1(1-\rho)(1+\rho) + \psi_2\hat{A}_t + \psi_2\hat{B}_t(1-\rho), \\ T(\hat{B}_t) &= (\psi_1\rho + \psi_2\hat{B}_t)\rho, \\ T(\hat{C}_t) &= \psi_1\rho + \psi_2\hat{B}_t, \end{aligned}$$

where the function $T : \mathbb{R}^N \to \mathbb{R}^N$ maps the estimated coefficients into the actual parameters, which are in turn determined by the estimates. There is a self-referential feature inherent in all learning models which can be seen in equation (2.27). Although the estimated parameters are non-stationary during their transition to REE values, learners neglect this fact, since a least squares method assumes the 'true' A, B and C to be constants. Intuitively, if the coefficient which determines the responsiveness to expectations is sufficiently small, then this specification error becomes asymptotically negligible and the economy converges to the REE (Evans and Honkapohja, 2001). The T-mapping to $\hat{C}_t$ is determined by the other coefficients, and the estimate $\hat{C}_t$ is independent of C and does not influence stability results. Therefore, in what follows we refer to the mappings $T\left(\hat{A}_t, \hat{B}_t\right)$ and for $\hat{C}_t$: $V(\hat{B}_t)$.

Let $z'_{t-1} = (1, y_{t-1})$, $x'_{t-1} = (\hat{A}_t, \hat{B}_t)$ and

$$\theta_t = z'_{t-1}x_{t-1} + \eta_t. \tag{2.28}$$

The estimation error η_t is perceived by the agents to be independently and identically distributed *iid*. However, due to the self-referential nature of the model there is an endogeneity bias which agents are unaware of, and thus η_t is *not* truly *iid*. We define $R_t = t^{-1}\sum_{i=1}^{t} z_{i-1}z'_{i-1}$, which allows us to write the RLS estimator as

$$R_t = R_{t-1} + t^{-1}(z_{t-1}z'_{t-1} - R_{t-1}), \tag{2.29}$$

$$x_t = x_{t-1} + t^{-1}R_t^{-1}z_{t-1}(\theta_t - z'_{t-1}x_{t-1}), \tag{2.30}$$

and thus

$$x_t = x_{t-1} + t^{-1} R_t^{-1} z_{t-1} \left(z_{t-1}' \left[T\left(\hat{A}_t, \hat{B}_t\right) - x_{t-1} \right] + V(\hat{B}_t)\varepsilon_t \right), \tag{2.31}$$

with the gain sequence $1/t$, often referred to as decreasing gain learning.[12] This gain guarantees that asymptotically new information is disregarded by agents.

The stability of the system in (2.29) and (2.31) with decreasing gain is governed by the following ordinary differential equation (ODE), where τ denotes 'notional' time:

$$\frac{d}{d\tau}\left(\hat{A}, \hat{B}\right) = T\left(\hat{A}, \hat{B}\right) - \left(\hat{A}, \hat{B}\right). \tag{2.32}$$

The REE is E-stable if (2.32) is asymptotically locally stable under learning (Evans and Honkapohja, 2001). This is the case, if all the eigenvalues of the Jacobian of $T(\hat{A}, \hat{B}) - (\hat{A}, \hat{B})$ have negative real parts. Here the necessary condition for E-stability is $\psi_2 \rho < 1$, with sufficient condition

$$\psi_2 = \delta\left[1 - \left(\lambda + \frac{\beta\bar{\theta}q(\bar{\theta})}{\alpha}\right)\right] < 1. \tag{2.33}$$

This holds for all possible well-defined sets of parameter values, and there is also global convergence to the REE (see B.1.4): $\delta \in [0,1),\ \lambda \in [0,1],\ \beta \in [0,1],\ \mu > 0,\ \alpha \in (0,1),\ c > 0$, and which all imply $\bar{\theta} \geq 0$. As explained in the previous section, the model is determinate if $|\psi_2| < 1$. We can therefore state the following:

Proposition 2.4.1 *If the economy described by the system* (2.21) *and* (2.22) *exhibits determinacy and the PLM is of the MSV form, and if agents learn using least squares updating, then so long as* $\psi_2 < 1$ *the unique REE is E-stable.*

In other words, the textbook linearised model of labour market search and matching frictions, with homogeneous agents and no-on-the-job search (Pissarides, 2000: Chapter 1), has a unique E-stable equilibrium. Sets of parameter values which move ψ_2 closer to one will imply slower convergence to the REE. It is intuitive and clear from (2.33) that these will be parameters which lessen the magnitude of the dynamics in the labour market, such as a small separation probability or low worker bargaining power.

[12] In the case of constant gain learning the weight given each observation is geometrically declining with the time since it was observed, and the gain sequence would be $0 < \gamma < 1$.

2.4.2 Strong E-stability of the MSV solution

One potential criticism of the econometric learning literature is that it is not clear how agents could settle upon a particular law of motion for the economy. Strong E-stability of a system is defined if the previous result is robust to over-parametrisation of the PLM (Evans and Honkapohja, 2001). Assume instead that agents are forming their expectations of θ_{t+1} according to the general ARMA representation (B.4), and are not endowed with a PLM of the MSV form. Moreover, due to econometric considerations they start with an arbitrarily over-parametrised version,

$$\theta_t = a + \sum_{j=1}^{s} b_j y_{t-j} + \sum_{j=1}^{r} c_j \theta_{t-j} + \sum_{j=1}^{q} d_j \varepsilon_{t-j} + \sum_{j=1}^{l} f_j \eta_{t-j} + d_0 \varepsilon_t + f_0 \eta_t. \quad (2.34)$$

Accordingly, expectations of θ_{t+1} take the form:

$$\theta^e_{t+1} = a + \sum_{j=1}^{s} b_j y_{t+1-j} + \sum_{j=1}^{r} c_j \theta_{t+1-j} + \sum_{j=1}^{q} d_j \varepsilon_{t+1-j} + \sum_{j=1}^{l} f_j \eta_{t+1-j}, \quad (2.35)$$

which can be substituted into equation (B.3) to obtain the new ALM and a corresponding T-mapping in the same way as before (see B.1.5). Let $\mathbf{b}' = (b_1, ..., b_s)$, $\mathbf{c}' = (c_1, ..., c_r)$, $\mathbf{d}' = (d_0, ..., d_q)$, and also $\mathbf{f}' = (f_0, ..., f_l)$. Further, define $\phi' = (a, \mathbf{b}', \mathbf{c}', \mathbf{d}', \mathbf{f}')$. According to the E-stability principle, the ODE governing the stability of the above system is given by

$$\frac{d\phi}{d\tau} = T(\phi) - \phi. \quad (2.36)$$

To investigate whether agents will detect the over-parametrisation and converge towards the MSV solution, the stability of (2.36) at the REE must be studied. In B.1.5 we show the following:

Proposition 2.4.2 *If the economy described by the system* (2.21) *and* (2.22) *exhibits determinacy and the PLM is of the over-parametrised ARMA form, and if agents learn using least squares updating, then so long as* $\psi_2 < 1$ *the unique REE is Strongly E-stable.*

2.5 Analysis

We present a brief analysis of the unemployment model with econometric learning described above. We consider two illustrative simulations to demonstrate the implied speed of convergence and dynamics of the model. First, we demonstrate E-stability when starting 'realistically' far away from the REE. Second, with agents initially assumed to have learned the REE, we consider the impact of a structural shift implied by an arbitrary change in some parameter value. We then discuss the speed of convergence and results with constant gain learning. We also consider the implications if we relax an implicit assumption that agents understand the joint dynamics of unemployment and vacancy creation.

2.5.1 Simulations

We follow an illustrative parametrisation strategy, using seasonally adjusted UK quarterly[13] data for the period 2002-2013 (see B.2 for a brief discussion of this strategy).[14] Table 2.1 gives the complete list of parameters and implied values of the endogenous variables for the deterministic steady-state equilibrium. Summary statistics of some UK labour market variables are described in Table 2.2, which are consistent with the parametrisation here.

TABLE 2.1: Assumed/estimated parameter values and steady-state equilibrium

Parameter	Assumed value
y - labour productivity	1
b - non-employment flow value	0.8
c - vacancy flow cost	0.25
λ - separation rate	0.023
μ - matching efficiency	0.56
α - matching elasticity	0.67
β - worker bargaining power	0.67
δ - discount factor	0.99
ρ - persistence of y	0.84
σ - std dev. of innovations to y	0.006
Endogenous variable	**Steady-state eq. value**
θ - tightness	0.35
u - unemployment	0.055
v - vacancy rate	0.019
w - wage	0.99

Source: authors' calculations.

[13] As pointed out by a referee, the timing structure of the model implies an average time between hiring and production for workers of one and a half months, and when calibrating a model with labour market search frictions it would be more generally preferable to use a monthly periodicity. But when we wish to capture the role of aggregate uncertainty affecting agents' decisions, since UK National Statistics are generally released quarterly, we believe our timing is justified. What matters for plausibly estimating the role of learning dynamics is the frequency at which it is assumed new aggregate data becomes available to the agents, since between times the agents' model parameters will remain unchanged.

[14] All data used and described are from the Office for National Statistics, accessed 01/08/2014. Labour market data are for those aged 16 and over. For a more complete calibration of the unemployment model using UK data see Burgess and Turon (2010).

For completeness we write out in full the stochastic recursive sequence that represents the adaptive learning model, stating from an initial period t_0:

$$
\begin{aligned}
&(I) \qquad u_{t+2} = \lambda(1-u_{t+1}) + \left[1-\mu\left(z_t'T(x_t)+V(x_t)\varepsilon_{t+1}\right)^{1-\alpha}\right]u_{t+1},\\
&(II) \qquad R_{t+1} = R_t + \frac{1}{t+1}(z_t z_t' - R_t),\\
&(III) \qquad x_{t+1} = x_t + \frac{1}{t+1}R_{t+1}^{-1}z_t\left\{z_t'[T(x_t)-x_t]+V(x_t)\varepsilon_{t+1}\right\},\\
&(IV) \qquad y_{t+1} = (1-\rho)+\rho y_t + \varepsilon_{t+1},\\
&(V) \qquad \varepsilon_{t+2} \sim i.i.d.\, N(0,\sigma^2).
\end{aligned}
$$

When written out in sequence order, the simultaneity which requires us to exclude y_t from the information set used to estimate x_t becomes clearer. The adaptive learning process, which takes place at the beginning of each period, can also be described by Figure 2.1. To initiate the sequence from t_0 we must choose initial values u_1,

FIGURE 2.1: Timeline of the labour market and agents' learning

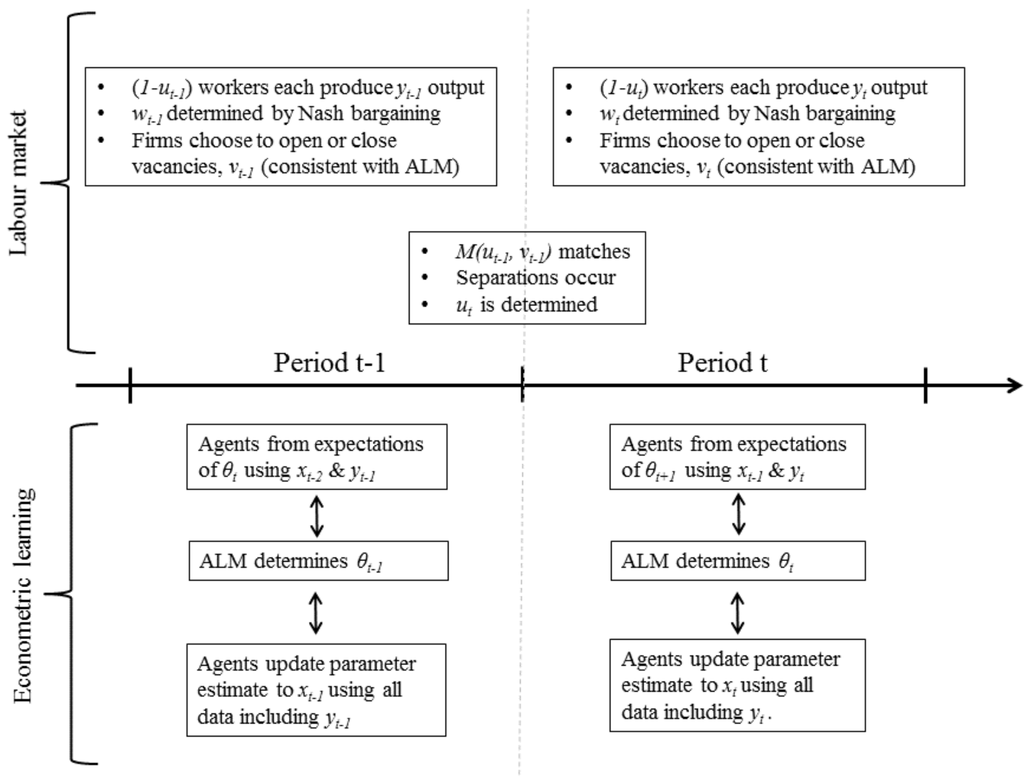

x_0, z_0 R_0 and ε_1. The asymptotic properties of decreasing or constant gain least squares recursion will hold irrespective of the initial conditions. As suggested by Carceles-Poveda and Giannitsarou (2007), the approach to setting initial values z_0 and

TABLE 2.2: Summary statistics of labour market states & quarterly transition rates: consistent with the model's parametrisation, 2002q1-13q2

	Mean	Std err.
Tightness - $\theta_t = \frac{v_t}{u_t}$	0.35	0.022
Job finding rate - $\theta_t q(\theta_t)$	0.39	0.011
Job separation rate - λ_t	0.023	0.00093
'Steady-state unemployment rate' - $u_t^* = \frac{\lambda_t}{\lambda_t+\theta_t q(\theta_t)}$	0.056	0.0041
Unemployment rate	0.057	0.0022

Source: authors' calculations using UK Labour Force Survey and Labour Market Statistics. The unemployment rate is the share of the economically active population ILO unemployed. The job finding and separation rates are consistent with in reality a three-state system, which includes inactivity; i.e. the job separation rate is not only the direct flow rate from employment to unemployment but in addition the indirect flow via inactivity (Smith, 2011). See B.2 for more details.

R_0 should depend on the particular model in question and the empirical purpose of the researchers. One approach could be to use historic or randomly generated data, with t_0 set sufficiently large such that R_0 is invertible; in this case $t_0 \geq 2$. This would be most appropriate when comparing the performance of models which assume that agents are 'good econometricians' against real data. However, this gives few clues as to how large t_0 should be, and the subsequent simulation is likely to be sensitive to this assumed level of agents' memory, particularly for decreasing gain least squares. Another attractive option is to choose initial values from an assumed distribution around the REE.

To set initial conditions here, using the same data used to parametrise the model, we estimate using least squares

$$\theta_t = \kappa_0 + cubtr_t + By_{t-1} + \kappa_1\zeta_{t-1} + \kappa_2\zeta_{t-2} + \zeta_t, \qquad t = 2001q3...2013q3, \quad (2.37)$$

where $cubtr_t$ represents a cubic time trend to address the possibility that agents could recognise low frequency structural breaks in the relationship, output per worker is normalised but not de-trended, and we include significant MA terms, to account for auto-correlation when the MSV is applied to real world data, which the good econometrician may in practice account for by what we have referred to before as an over-parametrised PLM.[15] Given an estimate of $\hat{B}$ from (2.37), we choose an initial value for $\hat{A}$ such that the economy is initially at $\bar{\theta}$, the deterministic steady-state equilibrium. $t_0 = 49$ is the maximum number of UK observations available. Using

[15]In determining initial conditions, one could also consider the class of GARCH, error correction, or even VAR models, however we believe this would be an unnecessarily significant leap from the straightforward least squares updating we assume that a 'good econometrician' carries out in practice, and which constitutes the learning algorithm we study here.

this approach, we set $R_0 = \left(\begin{smallmatrix} 1 & 1 \\ 1 & 1.0014 \end{smallmatrix}\right)$, $x_0' = (-1.42, 1.77)$, $z_0' = (1,1)$, $\varepsilon_1 = 0$ and

$$u_1 = \frac{\lambda}{\lambda + \mu \left(z_0' T(x_0)\right)^{1-\alpha}} \quad (= 0.265).$$

To analyse the impact of adaptive learning we focus on the simulated time paths of wages and the tightness parameter, which are independent of the choice of u_1. With the parametrisation described above, the REE parameters of the MSV solution are given by $x_{REE}' = (-0.70, 1.055)$. The elasticity of θ to productivity at the long-run average level is then around three, which is significantly lower than observed in the data.

Figure 2.2 demonstrates a simulation over a hundred quarters of wages, unemployment and labour market tightness for the baseline case of agents with RE.[16] Unsurprisingly, as is common with this class of models, and as described in Table 2.3 when compared with Table 2.2, the generated sample path under the REE significantly underestimates the variance of tightness in the UK labour market; i.e. the model does not generate a realistic magnitude of unemployment fluctuations over the business cycle, with the standard deviation being approximately a quarter of that observed in the data.

TABLE 2.3: Simulation results under the REE, decreasing and constant gain learning

Number of qtrs after init. val.	*20*	*100*		
	Std dev.	Std dev.	Min.	Max.
REE				
w	0.011	0.0085	0.97	1.01
u	0.00054	0.00053	0.054	0.056
θ	0.0015	0.0012	0.32	0.37
Decreasing gain				
w	0.012	0.0092	0.97	1.01
u	0.00072	0.00070	0.054	0.057
θ	0.0020	0.0016	0.31	0.38
Constant gain ($\gamma = 0.05$)				
w	0.012	0.0090	0.97	1.01
u	0.00073	0.00069	0.054	0.057
θ	0.0021	0.0016	0.31	0.38

Source: authors' calculations.

[16] See Appendix Figure B1 for the simulated paths of output per worker and shocks used in all simulations here. These were generated using the random number seed 42 in the Python Numpy application.

Figure 2.2 also shows the equivalent simulation results when agents learn the REE with decreasing gain, with initial estimates of the PLM parameters as described above. Figure 2.3 shows the path of these parameter estimates as agents learn from their forecast errors. The key result is that convergence is very slow, when agents are given a relatively small amount of historical data (12.5 years) and with initial estimates of the model parameters not unrealistically far from the true REE values. As shown in Appendix Figure B3, this takes thousands of years despite being exponential. This indicates that under adaptive learning, an economy could be persistently away from its REE level of unemployment, on the high or low side, even though agents are behaving rationally in the limited sense prescribed by the ‘good econometrician.’ In this sense, RE can be a poor approximation in terms of levels to a model with learning. One recommendation from this result is that when calibrating the Mortensen-Pissarides model, targeting second moments of the data should always be preferable, whereas not exactly hitting levels of the endogenous variables may not be too concerning.

As a further example, in Appendix Figure B4 and Figure B5 we simulate the model with no memory, and allow the agents to have guessed the correct initial parameter estimates, $x_0 = x_{REE}$, but suppose that there is an immediate negative twenty percent shock to the flow value of unemployment b. In the REE, due to the rise in the surplus of a match, firms immediately open more vacancies, and the unemployment rate falls. Under learning, the initial increase in θ is smaller. Therefore, unemployment falls more slowly as agents attempt to disentangle the effects of the structural shock from the stochastic process. In this sense, the response to the shock under learning leads to a less volatile path for unemployment. If actual labour market data contain the effects of frequent structural shocks of this kind, then econometric learning will not improve the ability of the standard search model to match their cyclical properties.

FIGURE 2.2: Simulations of the labour market model and agents' parameter estimates: a comparison of the REE, decreasing gain and constant gain learning

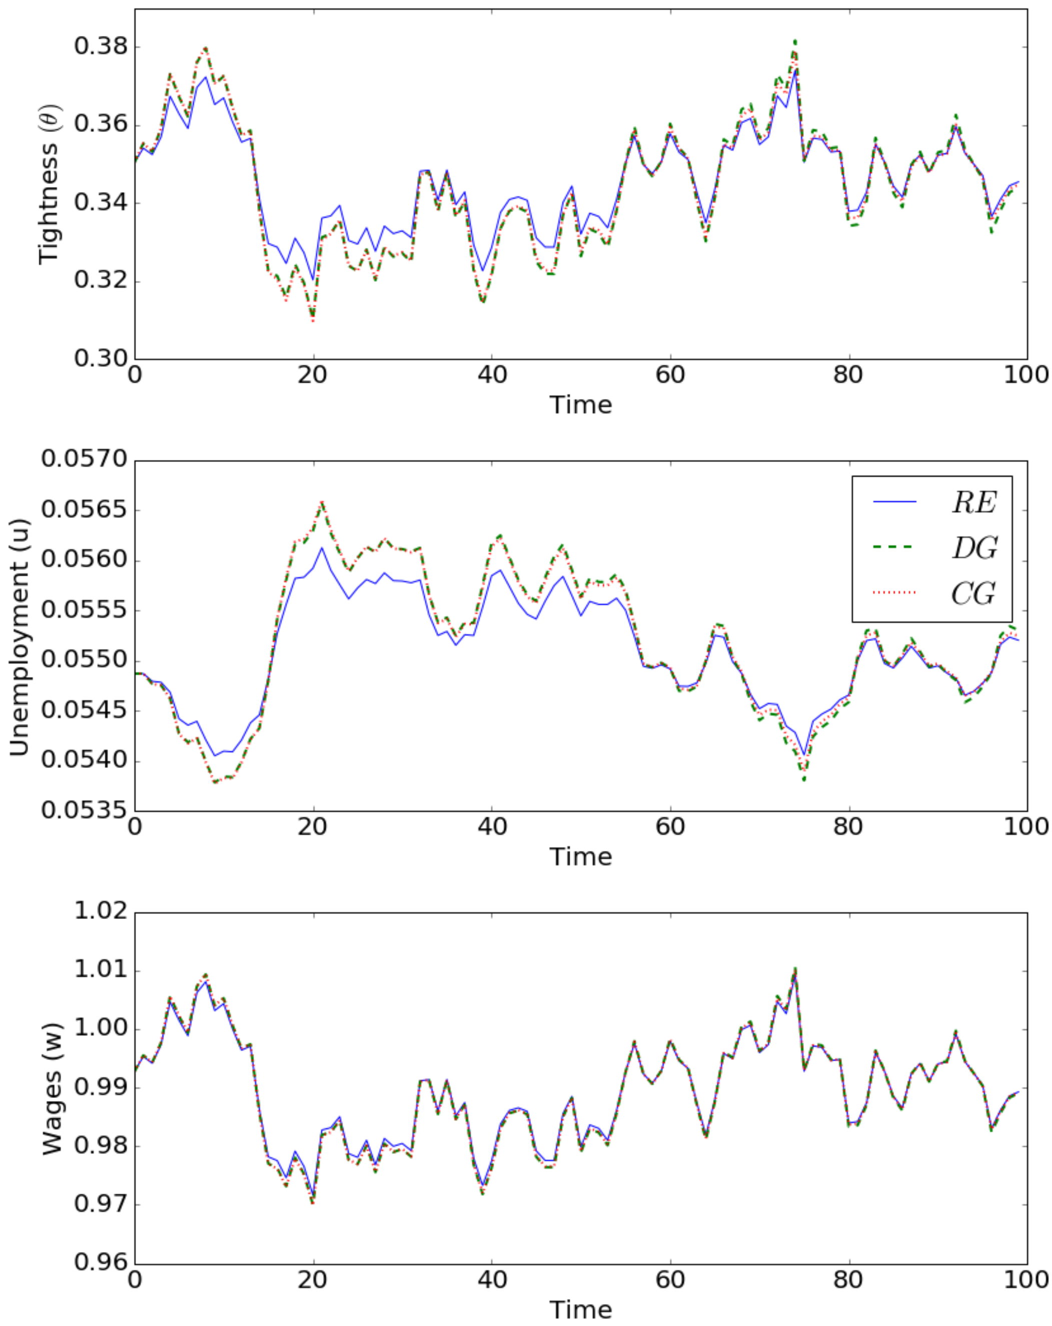

Note.- initial parameter estimates of the PLM, and $\hat{B}_0$ is assumed to be 'realistically' far away from the true REE values, whereas $\hat{A}_0$ is chosen such that under learning the economy begins at $\bar{\theta}$.

2.5.2 Speed of convergence

As shown theoretically in Benveniste et al. (2012), the learning of the agents results in root-t convergence to the true REE parameter estimates if all the eigenvalues of

FIGURE 2.3: Simulations of the labour market model and agents' parameter estimates: a comparison of the REE, decreasing gain and constant gain learning

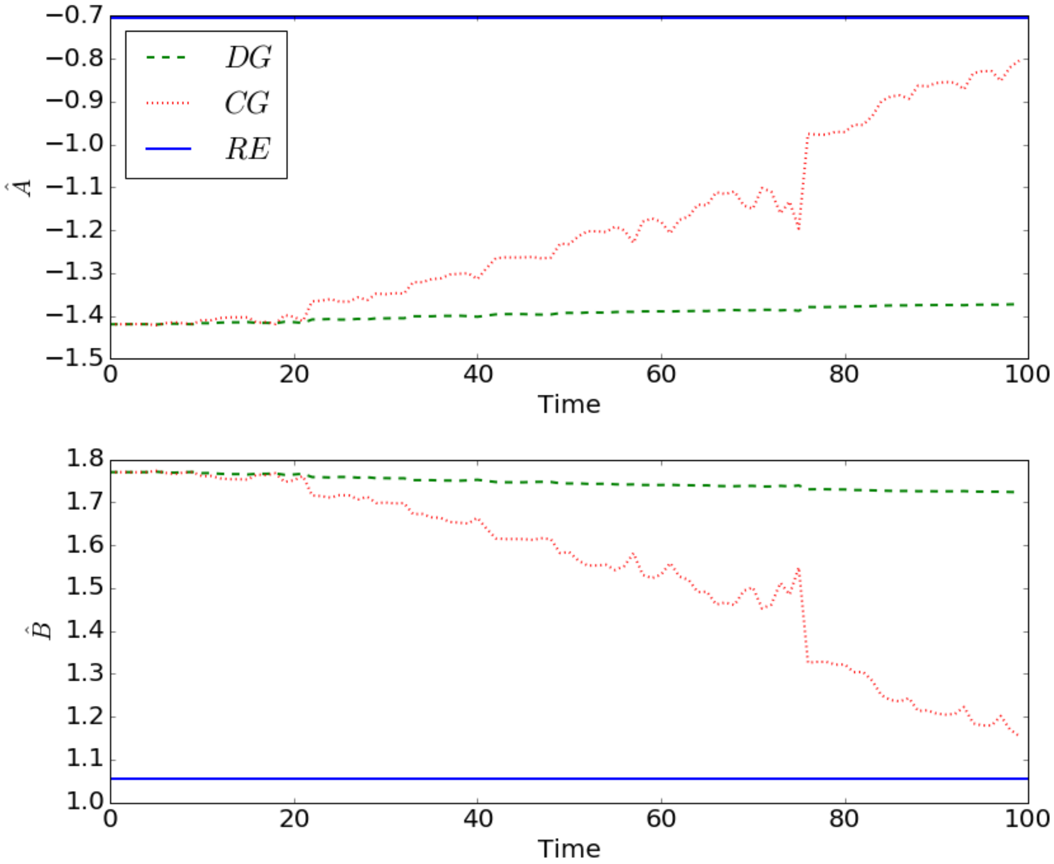

Note.- initial parameter estimates of the PLM, and $\hat{B}_0$ is assumed to be 'realistically' far away from the true REE values, whereas $\hat{A}_0$ is chosen such that under learning the economy begins at $\bar{\theta}$.

the system's Jacobian have a real part strictly less than a half.[17] Here this requires $\psi_2 < 1/2$. In the example parametrisation above this is not ensured, with $\psi_2 = 0.57$. More generally, it can be shown with simulations that the speed of convergence decreases substantially as $\psi_2 \to 1$, the threshold for E-Stability. To illustrate a decrease in the speed of convergence, in Appendix Figure B6 we consider a value of $\psi = 0.91$ by decreasing worker bargaining power to $\beta = 0.1$, keeping all other parameters except c constant, which is always used to match the mean value of θ from the UK data. As expected, the rate of convergence decreases, and the economy remains more persistently away from the REE. As such, choosing parameter values which guarantee a higher speed of convergence is one way in which the REE model could become an improved approximation of an alternative with econometric learning.

[17] I.e. the rate of convergence at which in classical econometrics the mean of the least squares parameter estimate converges to the true value.

2.5.3 Constant gain learning

In Figure 2.2 we also compare the results of our first simulation with decreasing gain learning to an equivalent example with constant gain parameter $\gamma = 0.05$.[18] When agents weight recent data more, convergence to the REE is faster, and agents' parameter estimates are more volatile. This faster convergence results in more volatile series of labour market tightness, wages and unemployment. However, the gain parameter generating this faster convergence roughly implies that agents only use data over the past twenty quarters to update their beliefs, and is notably outside the range suggested by the adaptive learning literature (see Di Pace et al. (2016) for a discussion). Our simulation results with constant gain learning and more reasonable levels of memory weighting are not dissimilar to those obtained with decreasing gain.

2.5.4 An alternative non-steady-state perceived law of motion

So far we have described a model of econometric learning in which agents endeavour to forecast labour market tightness θ. However this is a construct of the model and its assumptions. It is an attractive feature of the search and matching models that the equilibrium can be described by this single choice variable, determined by the state of the productivity process, but independent of unemployment. But firms in the model are described as choosing the number of vacancies to post, or analogously whether or not to enter the labour market. And for given levels of θ and productivity, this choice does depend on the state of the labour market. In the REE, if we consider the economy as initially being at some steady state (i.e. unemployment and vacancy rates are on the Beveridge curve), then in moving to any new steady state the vacancy rate changes non-monotonically. In characterising agents as learning how to choose and forecast θ, we imply that they fully understand the non-steady-state dynamics of the model. Here we consider the implications of relaxing this assumption (for what follows, see B.1.6 for complete derivations and descriptions of parameter values).

Linearising (2.14), (2.20) and (2.12) around the steady-state deterministic values of vacancies, employment and output per worker, we derive an alternative system defining the economy:

$$v_t = \kappa_0 + \kappa_1 y^e_{t+1} + \kappa_2 v^e_{t+1} + \kappa_3 n^e_{t+1} + \kappa_4 n_t, \tag{2.38}$$

$$n_t = \phi_0 + \phi_1 n_{t-1} + \phi_2 v_{t-1}, \tag{2.39}$$

$$y_t = (1-\rho) + \rho y_{t-1} + \varepsilon_t. \tag{2.40}$$

[18]For constant gain learning there is no analytic solution for expectational stability and so we must select a reasonably small gain parameter to ensure convergence.

We endow agents with a PLM in which they use both the output per worker and employment states to forecast vacancy creation,

$$v_t = \hat{A}_t + \hat{B}_t y_t + \hat{C}_t n_{t-1}. \tag{2.41}$$

Given (2.38)-(2.41), we can then derive the ALM for this version of the economy, and subsequently a $\tilde{T}$-mapping

$$\tilde{T}\left(\hat{A}_t, \hat{B}_t, \hat{C}_t\right) = \left(\tilde{\kappa}_0 + \tilde{\kappa}_2\left[\hat{A}_t + (1-\rho)\hat{B}_t\right], \tilde{\kappa}_1 + \tilde{\kappa}_2\rho\hat{B}_t, \tilde{\kappa}_3 + \tilde{\kappa}_2\hat{C}_t\right). \tag{2.42}$$

Assuming agents update their parameter estimates for the PLM using RLS as previously, and applying the same E-stability principle, it can be shown that the sufficient condition to guarantee local convergence to the REE is given by $\tilde{\kappa}_2 < 1$. Comparing this with the condition for stability of the MSV-PLM, $\tilde{\kappa}_2 \geq \psi_2$. Hence, convergence to the REE is slower when agents do not implicitly know the out of steady-state dynamics of employment and vacancy creation. The REE model is then a poorer approximation to an economy with econometric learning. What is more, for a subset of parameter values we cannot claim that the model is E-stable. For example, it is less likely to be E-stable in the circumstance of inefficiently high vacancy creation, departing from the Hosios (1990) condition (i.e. $\alpha > \beta$). Though for the parametrisation we have used here the model would still certainly converge to the REE.

2.6 Conclusion

We take the textbook linearised RBC version of the model of search and matching frictions for the labour market and show that the unique REE is not only always E-stable, for all well-defined sets of parameter values, but this result is robust to over-parametrisation of the MSV-PLM used by agents (Strong E-stability) with decreasing gain learning. These local convergence conditions also extend trivially to global convergence. Because the economy will eventually move to the REE when agents use econometric learning, the potentially unrealistic RE assumption in this class of model is nonetheless reasonable. We use recent UK data to parametrise the model, and show that although the model is E-stable, implied convergence can be very slow. Therefore, the RE model of unemployment fluctuations could in fact be a poor approximation to an economy in which agents more realistically learn as econometricians, especially in the presence of frequent structural or permanent shocks. The MSV-PLM implicitly assumes that agents understand the out of steady-state paths of employment and vacancy creation in the model. When we consider a version of the

PLM which relaxes this assumption, we see that convergence is further slowed, and local E-stability of the model is not guaranteed, making the approximation of the RE model even weaker.

Chapter 3

Measuring sectoral income shares: Accounting for input-output structures across countries

3.1 Introduction

Sectoral labour income shares provide a link from macroeconomic performance to households' perceptions of fairness and economic well-being (Atkinson, 2009), which makes labour income shares a primary concern for political economists and policy makers. Additionally, the recent surge of economic analysis in the fields of structural transformation and development accounting increases the need for data on labour income shares at the sectoral level. These shares do not only matter for quantitative exercises, but also qualitatively: differences between sectoral labour income shares within countries can drive structural transformation (Acemoglu and Guerrieri, 2008).

Despite their importance, comparable cross-country data on sectoral labour income shares do not exist.[1] Therefore, I compute labour income shares of the goods and the services sector for a large cross-section of countries. I find that these sectoral labour income shares significantly increase with the level of development. On average, the goods sector labour income share increases by 4.6 percentage points, and the services sector labour income share by 3.9 percentage points when an economy's output per person doubles.

[1]Evidence is limited to certain industries, for example, the network industries in OECD countries (Azmat et al., 2012), or manufacturing industries (Mareek and Orgiazzi, 2015). In developed countries, the latter typically accounts for only a small share of economic output.

Multi-sector models are usually formulated as value-added frameworks, in which output is produced using capital and labour as the only inputs. But industries also use intermediate inputs, and these inputs are commonly supplied by other industries. Because of these intermediate input linkages, constructing labour income shares at the sectoral level requires information on the input-output structure of an economy. Input-output tables from the World Input-Output Database (hereafter WIOD) provide all the necessary information for this exercise.

First, I compute the Total Requirements Matrix for each country and year in the WIOD for 1995-2009. Then I follow Valentinyi and Herrendorf (2008), and use these matrices to allocate capital and labour income across sectors, *as if* each sector used only its own output as intermediate input. In other words, I first derive value-added production functions for each country and category of final expenditure. Then I compute the goods sector labour income shares by aggregating all of those production functions which are associated with final expenditure on goods. The labour income share of the services sector follows similarly.

To adjust for the labour income of the self-employed, I first construct series of Net Operating Surplus for each industry, year, and country in the WIOD, using data collected from National Accounts. Then I use these series to impute the labour income of the self-employed, similar to Bernanke and Gürkaynak (2001) and Gollin (2002). Using different adjustment methods only affects the average level of the sectoral income shares, but does not notably change any other result in this paper. In contrast to Valentinyi and Herrendorf (2008), my estimates suggest that the labour income share of the services sector is on average slightly larger than the labour income share of the goods sector, and the differences between the goods sector and services sector labour income shares within countries are uncorrelated with the level of development across countries.

Further, I document that the labour income shares of the goods and the services sector increase significantly with the level of development, and hence the aggregate labour income share is also positively correlated with aggregate income. This contrasts with the results of Bernanke and Gürkaynak (2001) and Gollin (2002), who find that aggregate labour income shares are cross-sectionally uncorrelated with aggregate income. My results differ, because I use average labour income shares over 1995-2009, and thus I can exclude business cycle fluctuations as drivers of measured differences. Also, the sample of analysed countries differs from the above mentioned studies, mainly because the WIOD does not include African countries.

To assess the quantitative impact of using country-specific sectoral labour income shares vs. their corresponding U.S. values, I conduct a development accounting exercise. To allocate quantities of physical and human capital across the goods and the services sector, I use first-order conditions, which are derived from a simple cost minimisation problem of a representative firm. This approach is similar to Herrendorf and Valentinyi (2012), who also use first-order conditions to impute unobserved quantities of sectoral production factors, however, in contrast to their work, I do not assume that sectoral labour income shares are identical across countries. I find that the goods sector is less productive in poor countries than the services sector, supporting the findings of Herrendorf and Valentinyi (2012). Because sectoral labour income shares increase with development, using their corresponding U.S. values overestimates the elasticity of the production function with respect to labour input in relatively poor countries. I find that the sectoral productivity disparities are underestimated by nearly 20 per cent across countries.

3.2 Input-output tables: linking the production and the expenditure-side of the economy

To study economic phenomena and variables at the sectoral level, for example, structural change or productivity differences at the sectoral level across countries, economists have to make assumptions about sectoral production functions. Most commonly, these production functions are assumed to combine capital and labour inputs. Although production functions that do not restrict the elasticity of substitution between inputs to equal one are becoming increasingly popular, the Cobb-Douglas production function is still the benchmark in the literature (Herrendorf et al., 2014). For a given country, suppose that real output of sector $z \in \{G,S\}$, where G denotes the goods sector and S denotes the services sector, is produced according to:

$$y_z = A_z k_z^{1-\alpha_z} h_z^{\alpha_z} , \tag{3.1}$$

with total factor productivity A_z, the quantity of physical capital inputs k_z, and the quantity of human capital inputs h_z. The coefficient $\alpha_z \in (0,1)$ is the sectoral labour income share. As Valentinyi and Herrendorf (2008) note, this functional form abstracts from intermediate inputs. Formulating a sectoral production function without intermediate inputs implicitly assumes that each sector uses only intermediate inputs which it produced itself. However, the fact that large quantities of intermediate inputs are traded between the goods and the services sectors within countries (Grobovšek, 2017), means that abstracting from these intermediate input linkages is potentially not an innocuous assumption.

Taking the intermediate input linkages across sectors seriously requires data on the producers and consumers of these intermediate inputs - that is, one requires input-output tables. In the following explanation, I use the notation of the U.S. Bureau of Economic Analysis as outlined in Horowitz and Planting (2006).[2] Suppose the economy consists of n industries and m commodities, where industries combine intermediate inputs with capital and labour to produce output. Bold face lower case letters denote vectors and bold face upper case letters denote matrices throughout the paper.

Let $\mathbf{W}$ denote the Make matrix with dimension $(m \times n)$. The entries of a particular row of $\mathbf{W}$ show the value of the produced output of a particular commodity, which is supplied by each of the n industries, expressed as a share of the total value of the output of this commodity. For example, the commodity of "Motor Vehicles and Trailers" is produced mostly by the "Transport Equipment" industry, but is also produced by the "General Machinery" industry. For the latter, Motor Vehicles and Trailers is a secondary output, while for the Transport Equipment industry it is the primary output. Note that (3.1) implicitly abstracts from these secondary outputs. The $(m \times 1)$ vector $\mathbf{q}$ has the typical element $\{q\}_j$ which equals the total value of the output of commodity j.

Similarly, let $\mathbf{B}$ denote the Use matrix with dimension $(m \times n)$. The columns of the Use matrix $\mathbf{B}$ show the value of the commodity inputs required for the production of one unit of gross output of a given industry in U.S. Dollar. For example, the "Transport Vehicles" industry uses mostly the commodity of "Fabricated Metals", but also "Rubber and Plastic Products". Again, (3.1) abstracts from the use of these intermediate inputs, since they are not produced by the Transport Vehicles industry in this example. The $(n \times 1)$ vector $\mathbf{g}$ has the typical element $\{g\}_i$ which equals the total value of the output of industry i.

The k-th element of the $(m \times 1)$ vector $\mathbf{e}$ shows the value of the final expenditure on commodity k. These three vectors and matrices are related through accounting identities as follows:

$$\mathbf{q} = \mathbf{B}\mathbf{g} + \mathbf{e} \, , \tag{3.2}$$

$$\mathbf{g} = \mathbf{W}'\mathbf{q} \, . \tag{3.3}$$

[2]See also Miller and Blair (2009) for more details about the analysis of input-output data.

The first identity (3.2) states that the value of total domestic output of a commodity equals the value of this commodity used as intermediate input in the production of gross output by all industries plus the value of final expenditure on this commodity. The second identity (3.3) links the value of the gross output of an industry to the value of this industry's primary and secondary commodity output.

Denote the $(n \times n)$ identity matrix by $\mathbf{I}$. Combining the above identities (3.2)-(3.3) to eliminate the vector of commodity outputs $\mathbf{q}$ gives

$$\begin{aligned} \mathbf{g} &= \mathbf{W}'(\mathbf{I}-\mathbf{B}\mathbf{W}')^{-1}\mathbf{e} \, , \\ &:= \mathbf{L}\mathbf{e} \, , \end{aligned} \tag{3.4}$$

where the last line introduces the $(n \times m)$ Leontief matrix $\mathbf{L} = \mathbf{W}'(\mathbf{I}-\mathbf{B}\mathbf{W}')^{-1}$. This matrix is an industry-by-commodity Total Requirements Matrix, where entry $\{L\}_{ij}$ gives the value of industry i's gross output which is required to deliver one U.S. Dollar of commodity j to final use.

Let $\alpha_z \in (0,1)$ denote the labour income share of sector $z \in \{G,S\}$, and let $\mathbf{r}$ be an $(n \times 1)$ vector, with entry $\{r\}_i$ equal to the ratio of labour income to gross output in industry i. Similarly, let $\mathbf{V}$ be an $(n \times 1)$ vector, with entry $\{V\}_i$ equal to the ratio of value-added to gross output in industry i. With this notation, the labour income share in sector z is

$$\alpha_z = \frac{\mathbf{r}'\mathbf{L}\mathbf{e}_z}{\mathbf{V}'\mathbf{L}\mathbf{e}_z} \, . \tag{3.5}$$

The entries of expenditure vector $\mathbf{e}_z$ equal the final expenditure on commodity j if $j \in z$ and zero otherwise. The vector $\mathbf{r}'\mathbf{L}$ computes the value of labour income in each final expenditure category, and the vector $\mathbf{V}'\mathbf{L}$ computes value-added in each final expenditure category. Weighting both vectors by the vector $\mathbf{e}_z$ then computes the ratio of labour income to value-added in sector z. Appendix C.2 shows the final expenditure on commodities included in each sector.

3.3 The World Input Output Database and other data sources

I use data from the WIOD, 2013 Release (see Timmer et al. (2015) for details on the construction of this dataset).[3] The WIOD contains annual data covering 40 countries - 27 European Union members (Croatia being the one EU member state not available, as of November 2017) and 13 other major economies, including China, India, and

[3]Available at http://www.wiod.org/release13; accessed 01/09/15

the United States, for the period from 1995 to 2011. I exclude Mexico, Slovakia, and Turkey from the following analysis, because required additional data were not available to adjust for the labour income of the self-employed. Not excluding these three countries would likely strengthen the finding that sectoral labour income shares are increasing with the level of development.

The WIOD contains input-output tables as well as Socio Economic Accounts (SEA), which provide information about economic variables at the industry level, including hours worked by employees and self-employed, labour compensation, and value-added. The WIOD is constructed using only publicly available national account statistics and therefore can be adjusted and updated if new data becomes available or old data is revised.

I use the International Supply and Use tables of the WIOD, because these tables exclude imports from the intermediate inputs and final expenditures. I focus exclusively on the domestic economy and results thus reflect domestic labour income shares. To supplement the data in the WIOD, I collect data on Net Operating Surplus and Net Mixed Income (NOPS) for each industry, year, and country. Data on NOPS is obtained mostly from national accounts and from EUROSTAT, and the country sources are listed in Appendix C.1. The same data sources were used to construct the NOPS series and the WIOD, thus guaranteeing consistency between the two data sets. NOPS consists of two parts: profit of incorporated businesses after paying labour input costs and making allowances for the consumption of fixed capital, and Net Mixed Income, which is the equivalent measure for unincorporated businesses (small family businesses and self-employed). Therefore, NOPS provides an upper bound for labour income of the self-employed. As explained below, one would like to have data on Net Mixed Income instead of NOPS, but most countries do not provide data on these two variables separately at the industry level.

Additionally, I use data from the Penn World Table 9.0 (PWT): (1) expenditure-side real GDP, using prices for final goods that are constant across time and countries; and (2) the quantity of the capital stock using prices for structures and equipment that are constant across countries (Feenstra et al., 2015).[4] I also compute an aggregate index of human capital based on the average years of schooling from Barro and Lee (2013) which I use in the development accounting exercise below.

[4]Available at http://www.rug.nl/ggdc/productivity/pwt/; accessed 25/08/17.

3.3.1 Adjustments for the labour income of the self-employed

Total labour income in an industry is the sum of two components: employees' compensation and labour income of the self-employed. In the WIOD, wages and salaries of employees as well as other forms of remuneration in return for work, such as the contribution of employers to social insurances and pensions, are obtained from payroll data and labour force surveys for each country and industry. Additionally, information on the hours worked of employees from national labour cost surveys and enterprise surveys is used to compute the total compensation of employees.

However, as Krueger (1999) emphasises, national account statistics do *not* account for the second component, labour income of the self-employed, although some share of the reported earnings of the self-employed should clearly be attributed to labour income, with the reminder then comprising returns on invested capital, land rents, or monopoly profits. Different methods have been proposed in the literature to account for the labour income of the self-employed. These methods are: (1) considering only employees in the corporate sector; (2) assuming that the self-employed earn the same wage per hour as employees; (3) using additional data on operating surplus for each industry to impute profits of the self-employed; (4) applying a combination of methods (2) and (3) by assuming that self-employed earn the same wage per hour as employees, but this value can't exceed the total surplus of the self-employed.

To answer how the different adjustment methods of the labour income share affect its relationship with the level of economic development, I estimate the following regression using least squares:

$$\bar{\alpha}_i = a + \beta \; \overline{rGDPe_i} + u_i \; , \tag{3.6}$$

where $\bar{\alpha}_i$ is the average aggregate labour income share over 1995-2009 in country i, and $\overline{rGDPe_i}$ denotes the log of the geometric average of real GDP per person from the expenditure-side (from the PWT), in prices that are constant across time and countries (hereafter International Dollar), in country i over 1995-2009.[5] I compare averages to exclude business cycle effects. Figure 3.1 displays the labour income share estimates for each of the adjustment methods discussed in the following.

The first adjustment method simply excludes the self-employed, and considers only employees on payroll (Karabarbounis and Neiman, 2014). This method has the advantage of being transparent, and its results reflect payments to labour (Elsby

[5] Results presented here are robust to different measures of development: using real GDP per worker from the expenditure side, real GDP per person/worker from the production-side, or arithmetic instead of geometric averages does not change results notably.

FIGURE 3.1: Aggregate labour income shares across countries, 1995-2009: comparison of adjustment methods for the labour income of the self-employed

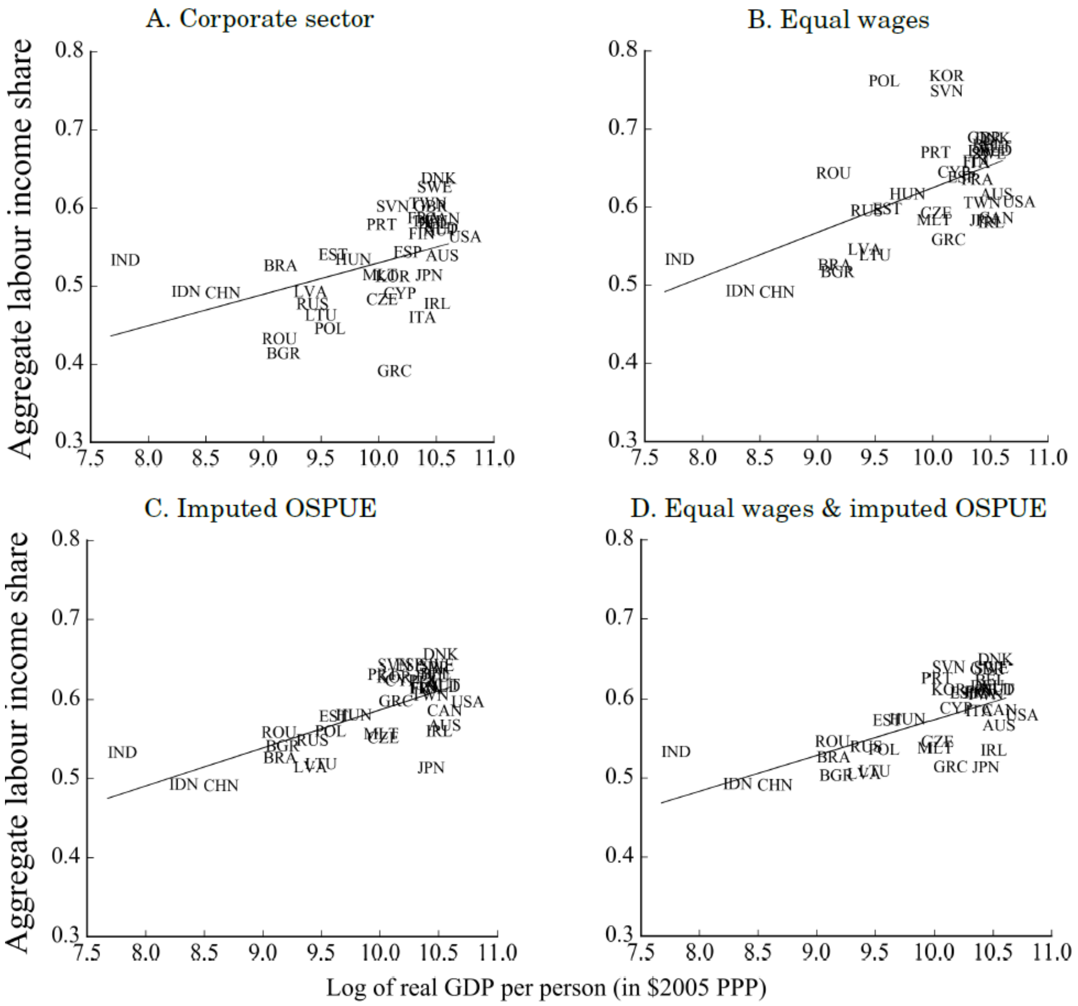

Notes.- plot of the average aggregate labour income share of each country against the geometric average of real GDP per person in International Dollar over 1995-2009. Solid line shows the estimated slope coefficient from (3.6). Luxembourg is excluded in figures, but is included in the regression.

et al., 2013) unambiguously. However, this assumption potentially excludes a large share of labour income. Especially in cross-country settings, where the number of self-employed varies substantially across countries, this approach seems problematic. For example, while the share of self-employed among all working individuals in the U.S. is on average 6.9 per cent over 1995-2009 according to the WIOD SEA, the corresponding values for Portugal and India are 19.2 and 55.4 per cent, respectively. Attributing only payroll compensation to labour income could systematically underestimate the share of labour income in poorer countries relative to rich countries. The first row of Table 3.1 shows that the share of payroll compensation to economy-wide value added is 52.2 per cent, averaged across all years and countries in the sample, and Figure 3.1A shows the payroll labour income share for each country.

The third column of Table 3.1 displays the coefficient estimates $\hat{\beta}$ for each adjustment method. The significant estimate of 0.04 for the payroll labour income share means that this share increases by 0.04 percentage points when real GDP per person increases by one per cent. Gollin (2002) also finds a positive relationship between the aggregate payroll labour income share and economic output per person. However, if the self-employed have a systematically higher labour income share in poorer countries than in richer countries, then the aggregate labour income share could actually be uncorrelated, or even negatively correlated, with the level of development.

TABLE 3.1: Comparison of different labour income adjustments, and regression estimates, 1995-2009

	Average 1995-2009 (1)	Standard deviation (2)	Slope estimate (3)	Spearman coefficient (4)
1. Only corporate sector	0.522	0.061	0.040^{***} (0.013)	0.81
2. Equal hourly wages	0.614	0.070	0.057^{***} (0.014)	0.88
3. Labour income equal to OSPUE	0.565	0.049	0.044^{***} (0.009)	0.95
4. Equal wages constrained by OSPUE	0.548	0.048	0.045^{***} (0.008)	1.00

Notes.- column (1) shows the average of labour income shares over 1995-2009. Column (2) shows the standard deviation of labour income shares over 1995-2009. Column (3) shows estimates $\hat{\beta}$, obtained from regression (3.6). Column (4) shows Spearman's rank correlation coefficient for pairwise comparison of the country-ranking of each row with the country-ranking of the fourth row. Newey-West standard error estimates in parenthesis.
*** $p<0.01$, two-sided test.

A second approach of accounting for labour income of the self-employed is to assume that the self-employed earn the same hourly wage rate as employees. This is the method which the U.S. Bureau of Labor Statistics (BLS) uses to compute their headline measure of labour income. An advantage of this method is that it accounts for differences in the hours worked between employees and self-employed. However, Gollin (2002) argues that the imputation of equal wages will be a poor approximation if systematic differences between the earnings of self-employed and employees exist. Because this assumption seems particularly unrealistic in emerging economies, the WIOD uses household surveys and census data to determine the labour income of the self-employed in relatively poor countries. These countries are: Brazil, China, India, Indonesia, Taiwan. The average labour income share increases to 61.4 per cent when this method is used (second row Table 3.1), and for some countries with a particularly high ratio of self-employed-to-employees (e.g., Poland and Korea) this method implies labour income shares larger than 75 per cent (Figure 3.1B).

Gollin (2002) proposes to take the operating surplus of private unincorporated enterprises (OSPUE) into account, arguing that most of this surplus will be labour income of the self-employed. However, because OSPUE data is not available for most of the WIOD countries, I follow the approach of Bernanke and Gürkaynak (2001): they argue that in countries with a relatively large share of self-employed, unincorporated surplus should also be relatively large compared to incorporated surplus. Hence, Bernanke and Gürkaynak assume that the ratio of incorporated to unincorporated surplus in total NOPS is identical to the ratio of employees to self-employed. This adjustment using *imputed* OSPUE has the advantage of being straightforward and transparent. However, if incorporated businesses generate higher returns per employee than unincorporated businesses, then splitting NOPS based on labour-input ratios will overestimate the share of unincorporated surplus in NOPS. Another disadvantage is that all the imputed OSPUE is assumed to be labour income, which is certainly an overestimation because, according to Gollin (2002), even in poor countries self-employed tend to use considerable quantities of capital inputs. The third row of Table 3.1 shows that the resulting average labour income share equals 56.5 per cent.

The last estimate of the self-employed labour income is computed by assuming that the self-employed earn the same hourly wage rate as employed workers (as in adjustment method (2)), but additionally I assume that the imputed OSPUE is the maximum of the labour income of the self-employed. This method has the advantage of using most of the available information by accounting for differences in hours worked between payroll employees and self-employed. However, especially in less developed countries where this assumption seems more unlikely to hold, imputed OSPUE is an upper bound for the labour income of the self-employed. As the fourth row of Table 3.1 shows, the average labour income share equals 54.8 per cent, which is by construction larger than the value of the payroll labour share, and smaller than for the other two estimates in the second and third row. Column (2) of the same table displays the standard deviation for each estimate. Because the standard deviation is the smallest for the estimate in the fourth row, and this estimate uses most of the available information, this is the preferred method. The associated coefficient estimate of (3.6) equals 0.045 and is significantly positive, as shown in the third column of Table 3.1.

To understand how the different estimates affect the relative ranking of countries based on their respective labour income share, I compute Spearman's rank correlation coefficient between the preferred estimate of the labour income share and all other estimates. This coefficient equals one if a perfect positive monotone relationship exists between the resulting rankings. For pairwise comparison of each adjustment method with the preferred adjustment method in the last row, the resulting Spearman coefficients are shown in the fourth column of Table 3.1. The rank correlation between adjustment methods which take the labour income of the self-employed into account is 0.88 and 0.95. This indicates that there is a strong positive monotone relationship between the different estimates of the average aggregate labour income shares, and so this ranking is relatively robust to different adjustment methods.

The aggregate labour income shares are increasing in the level of development, as measured by real GDP per person in International Dollar. This pattern is visible in all panels of Figure 3.1, and confirmed by the significantly positive slope estimates in Table 3.1. Most importantly, this result does not depend on the method used to impute the labour income of the self-employed. For the preferred estimate of the labour income share in the fourth row of Table 3.1, the last column shows that the labour income share significantly increases by 0.045 percentage points when real GDP per person increases by one per cent.

3.4 Main empirical results: Sectoral labour income shares across countries

Figure 3.2 plots the labour income shares of the aggregate economy, the goods sector, and the services sector across countries against the log of real GDP per person in International Dollar. Labour income shares are arithmetic averages over 1995-2009, and real GDP per person refers to the geometric average over the same period. Appendix Table C.1 and C.3 display the classifications of commodities into the goods and services sector and the sectoral labour income shares, respectively.

Figure 3.2 suggest that the labour income shares in the goods and the services sector increase with the level of development across countries. Table 3.2 displays the means and the estimated slope coefficients from regression (3.6), using as dependent variable the aggregate and sectoral labour income shares, and as explanatory variable the log of real level of expenditure-side GDP per person in International Dollar, all variables averaged over 1995-2009. The first row of Table 3.2 shows that the labour income share is slightly smaller in the goods sector than in the services sector. However, the hypothesis that the means of the sectoral labour income shares are identical can not be

FIGURE 3.2: Sectoral labour income shares, 1995-2009 averages

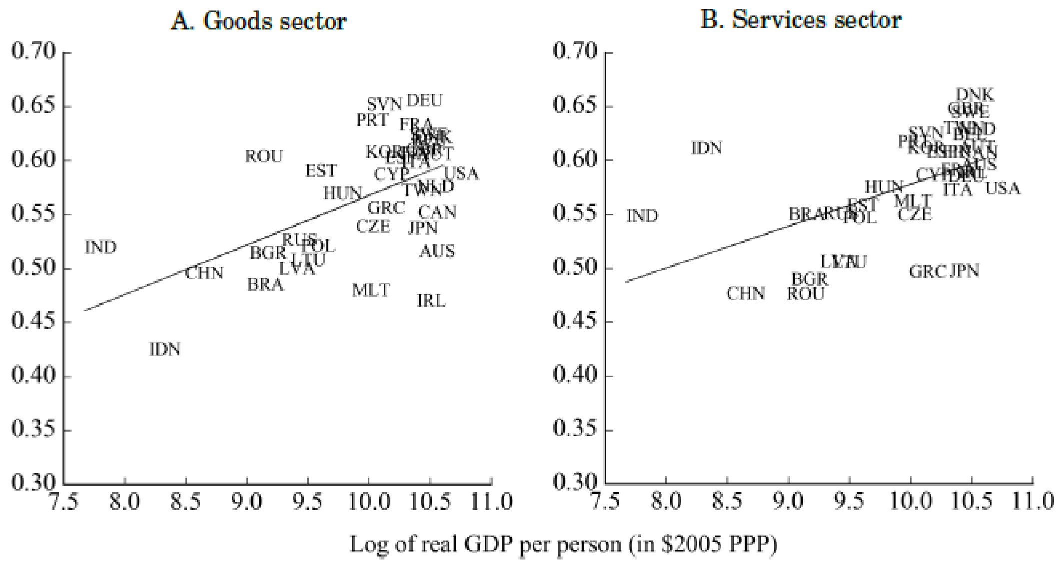

Notes.- goods sector and services sector labour income shares. Averages over 1995-2009. The horizontal axes show the geometric average over 1995-2009 of log real GDP per person from the expenditure-side in International Dollar (PWT). The solid line shows the least squares regression line from (3.6), varying the dependent variable accordingly. Luxembourg is omitted from this figure, but included in the regression.

rejected at the one per cent level: assuming independent sampling or paired sampling for the statistical test does not change this result.

TABLE 3.2: Means and coefficient estimates for sectoral labour income shares, 1995-2009

	Aggregate economy (1)	Goods sector (2)	Services sector (3)
1. Arithmetic mean	0.550	0.546	0.554
2. Estimate $\hat{\beta}$	0.045***	0.046***	0.039***
	(0.009)	(0.012)	(0.011)

Notes.- the first row shows arithmetic means computed over countries and years. The second row shows coefficient estimates from regression (3.6). Newey-West standard error estimates in parenthesis.
*** p<0.01, two-sided test.

The second row of Table 3.2 suggests that the sectoral labour income shares are increasing with the level of development: the coefficient estimates of the slopes are significant and positive at the one per cent level. When real GDP per person in International Dollar increases by one per cent, then the goods and services sector labour income shares increase by 0.046 and 0.039 percentage points, respectively. In other words, suppose that income per person doubles, then the expected increase in the labour income shares of the goods sector and the services sector is 4.6 and 3.9 percentage points respectively.

Figure 3.3 displays the differences between the goods and the services sector labour income shares in each country. These within-country differences and the level of development are not correlated. This is due to the previous finding that both the services and goods sector labour income shares increase with the level of development, thus keeping the difference approximately constant.

FIGURE 3.3: Within-country differences in labour income shares between the services and the goods sectors, 1995-2009 average

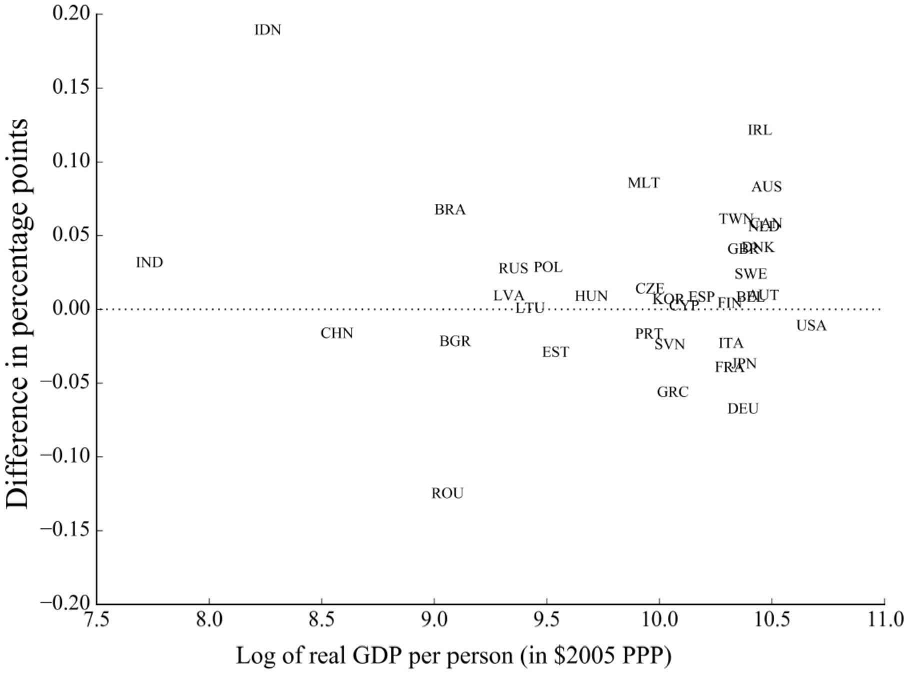

Notes.- horizontal axis shows geometric average of real GDP per person in International Dollar over the period from 1995-2009. Vertical axis shows the difference in labour income shares in percentage points, services sector less goods sector. Luxembourg is excluded in this figure.

Figure 3.3 also shows substantial differences within some countries: for example, the services labour income share in Indonesia exceeds the goods sector labour income share by nearly 20 percentage points. In contrast, the services sector labour income share in Romania is over 10 percentage points lower than in the goods sector. These relatively large differences are a significant finding, because they can be drivers of structural transformation (Acemoglu and Guerrieri, 2008). In their theoretical framework, the elasticity of substitution between the outputs of the goods and the services sector determines whether the sector with the relatively higher labour income share grows or declines in terms of capital and labour allocation. If the outputs are complements, countries positioned above the dotted line in Figure 3.3 will allocate a larger share of their capital and workforce to the goods sector as their economies grow, and *vice versa*.

The results in this section suggest that the labour income shares of the goods and the services sectors do not have the same value across the countries in the WIOD. Moreover, there is a significantly positive relationship between the level of development and the level of the sectoral labour income shares. This relationship is stronger for the goods than for the services sector. Thus, the assumption that sectoral labour income shares are equal to their respective U.S. values is not supported by the data. Both the levels of the labour income shares and the within-country differences are not constant within a relatively homogeneous group of mostly developed countries.

3.5 Decomposition of differences in labour income shares across countries

To better understand the causes of the observed differences in sectoral labour income shares across countries, this section decomposes the sectoral labour income shares. Since the usual practice in the literature is to assume that sectoral income shares are identical to their corresponding U.S. values, I choose the U.S. as a benchmark.

Denote the labour income share of sector z in country i by α_{iz}. The difference in the sectoral labour income share between country i and the U.S., in percentage points, is:

$$\alpha_{iz} - \alpha_{USz} = \frac{\mathbf{r}_i'\mathbf{L_i}\mathbf{e}_{iz}}{\mathbf{V}_i'\mathbf{L_i}\mathbf{e}_{iz}} - \frac{\mathbf{r}_{US}'\mathbf{L_{US}}\mathbf{e}_{USz}}{\mathbf{V}_{US}'\mathbf{L_{US}}\mathbf{e}_{USz}} , \tag{3.7}$$

where the notation is the same as in section 3.2, but here I make the country-dependency explicit. Production-side labour income shares are linked through the input-output structure to the expenditure-side, where they are weighted using the expenditure shares of final use. The observed difference is decomposed into these components as follows (see Appendix C for derivations):

$$\alpha_{iz} - \alpha_{USj} = \underbrace{(\Delta\mathbf{S}_{iz})'\mathbf{L}_i\mathbf{e}_{iz}}_{\text{VA labour share}} + \underbrace{\mathbf{S}_{iz}'(\Delta\mathbf{L}_i)\mathbf{e}_{iz}}_{\text{Supply chain}} + \underbrace{\mathbf{S}_{iz}'\mathbf{L}_{US}(\Delta\mathbf{e}_{iz})}_{\text{Expenditure weights}} . \tag{3.8}$$

The first term on the right-hand side represents differences in the production-side labour income share in value-added (hereafter VA labour share). For example, if country i had the same input-output structure and expenditure weights of final use as the U.S., then observed differences in expenditure-side labour income shares could only be caused by differences in the ratio of labour income to value-added at the production-side. The second term captures the effects of differences in the input-output network (supply chain) between the U.S. and country i. While not impacting on the aggregate labour income share, linking value-added labour income shares differently to the expenditure-side can result in varying sectoral labour income shares. For example, relatively high value-added labour income shares might be linked to the goods sector in the U.S., but to the services sector in country i. The last term represents difference in the weights which are used to aggregate the underlying final expenditure labour income shares into the sectoral labour income shares.

Figure 3.4 shows the contribution of each component to the observed difference in sectoral labour income shares between the U.S. and country i, separately for the goods and the services sectors.

The horizontal axes of each panel in Figure 3.4 show the observed differences in labour income shares, the left-hand side of (3.8). The vertical axes show the counterfactual differences, in percentage points, resulting from imposing on country i two of the three components that form the sectoral income shares of the U.S. For example, panel A imposes the U.S. supply chain and expenditure weights of the goods sector on country i, but keeps the observed value of the value-added labour income shares as in country i. The resulting differences between the sectoral labour income share in country i and the U.S., attributed to differences in the value-added labour share, are shown on the vertical axes. The U.S. value is centred at coordinate $(0,0)$ in all panels. The proximity to the 45 degree line indicates visually how much of the observed total differences between sectoral labour income shares each component explains.

Figure 3.4 suggests that the main drivers of the observed differences in labour income shares between the U.S. and other countries are differences in value-added labour income shares. Differences in the supply chains or the expenditure weights contribute relatively little to the observed differences in sectoral labour income shares. Table 3.3 displays the average contribution of the different components to the observed total average difference between sectoral labour income shares of the U.S. and all other countries.

The first row of Table 3.3 shows that the value-added labour income shares explains on average 2.8 percentage points (column (1)) of the observed average difference of 3.7 percentage points (column (4)) in labour income shares of the goods sector between the U.S. and all other countries. This suggests that approximately 75 per cent of cross-country variation in the goods sector labour income share originates from the production-side. The value of -0.3 percentage points in the first row of Table 3.3,

FIGURE 3.4: Decomposition of differences in expenditure-side labour income shares between the U.S. and other countries into value-added, supply chain, and expenditure weight, 1995-2009 averages

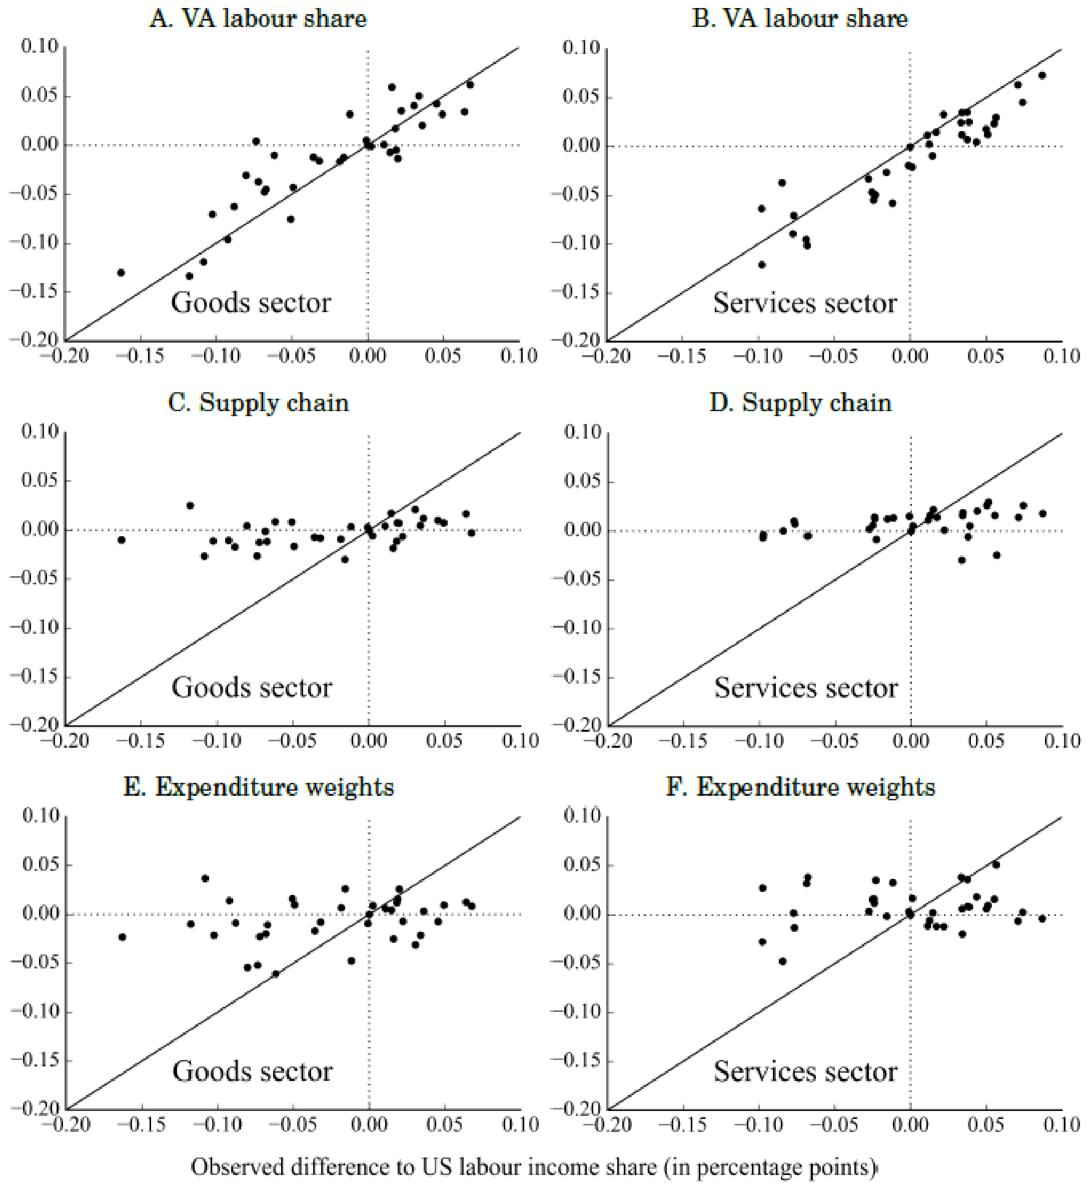

Notes.- decomposition of sectoral labour income shares according to equation (3.8), individual components averaged over the period 1995-2009. Each dot represents a country. The horizontal axes show the observed difference in percentage points to the U.S. sectoral labour income shares, $\alpha_{iz} - \alpha_{USz}$. Vertical axes show the differences to the U.S. labour income share in percentage points caused by the considered component, as indicated in the title of each panel. The U.S. value is centred at $(0,0)$.

column (2), suggests that the goods sector in other countries uses on average more intermediate inputs of industries with relatively low value-added labour income shares. Similarly, column (3) suggests that countries spend on average a larger share of their

TABLE 3.3: Decomposition of sectoral labour income shares

	VA labour share (1)	Supply chain (2)	Expenditure weights (3)	Total (4)
1. Goods Sector	−0.028	−0.003	−0.006	−0.037
2. Service Sector	−0.027	0.007	0.005	−0.015

Notes.- decomposition of expenditure side labour income shares according to equation (3.8), individual components averaged across countries and over the period 1995-2009. Column (4) shows the observed difference in percentage points to the U.S. sectoral labour income shares, $\alpha_{iz} - \alpha_{USz}$.

income, compared to the U.S., on goods-sector commodities which have relatively low labour income shares.

The second row of Table 3.3 shows the results for the decomposition of the total observed difference in labour income shares of the services sector between the U.S. and other countries. If all countries in the WIOD had supply chains and expenditure weights as in the U.S. (column (1)), then the services sector labour income shares would be on average 2.7 percentage points below the corresponding U.S. value, instead of only 1.5 percentage points. As column (3) shows, this is due to other countries consuming relatively more commodities of the services sector with relatively higher expenditure-side labour income share. Additionally, column (2) shows that these commodities are produced using relatively more intermediate inputs with higher value-added labour income shares.

The decomposition results suggest that the positive correlation between the level of development, as measured by real GDP per person, and sectoral labour income shares is related to the production-side labour income shares in general, and less associated with specific differences of the supply chain network or the structure of final expenditure across countries. While these two components tend to further decrease the labour income share of the goods sector, the decomposition suggests that they are increasing the labour income share of the services sector, relative to the corresponding U.S. value.

3.6 Approximations by the value-added labour income shares and the aggregate labour income shares

The data requirements for computing sectoral labour income shares, which take the full input-output structure into account, are relatively high; input-output tables are not always available, and classifications and accounting standards vary significantly across countries. Therefore, the finding that differences in sectoral labour income shares between countries are mainly caused by differences in production-side labour income shares suggests that these shares can provide a reasonable approximation to sectoral labour income shares.

Figure 3.5 plots the sectoral labour income shares against value-added weighted production-side labour income shares, separately for the goods and the services sectors, for each year and country. The 45 degree line provides a visual benchmark for a perfect

FIGURE 3.5: VA-weighted industry labour share approximation

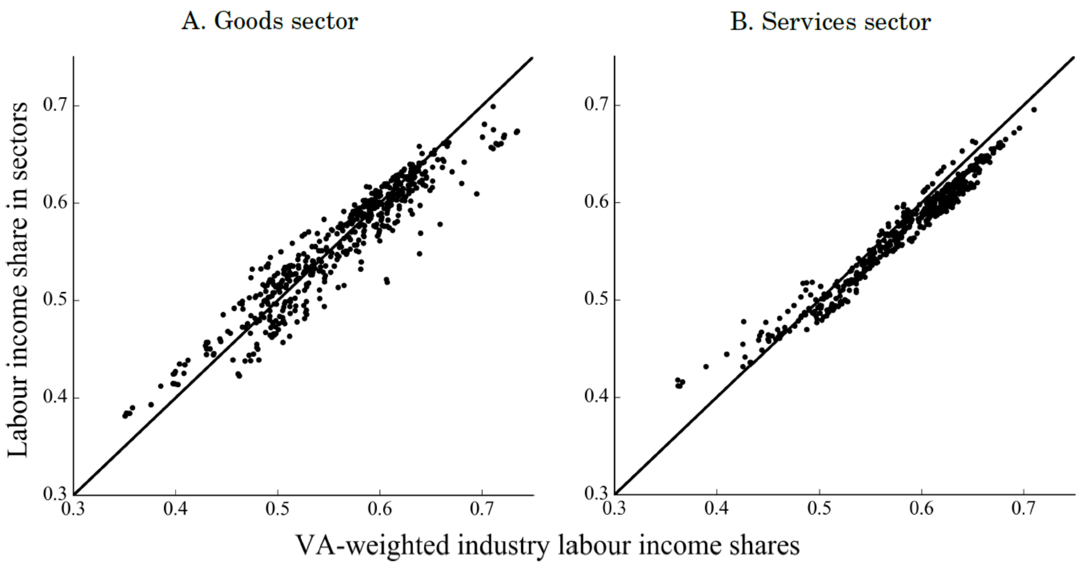

Notes.- value-added weighted industry labour income shares (horizontal axes) and sectoral labour income shares (vertical axes) for each year and country. The solid line represents 45 degrees.

fit. Value-added production-side shares are not systematically different from sectoral labour income shares, as expected from the decomposition results. The fit is better for the service sector than for the goods sector, but both approximations are significantly better than assuming the corresponding sectoral labour income shares of the U.S. apply across countries.

Although using value-added labour income shares does not require knowledge of the input-output structure, these income shares are also not available for a large number of countries. Therefore, I provide a visual assessment of the goodness-of-fit of the approximation by the aggregate labour income share, because these aggregate shares are available for almost all countries. Results are shown in Figure 3.6.

The interpretation is similar to the one of Figure 3.5: the 45 degree line indicates exact matches between aggregate and sectoral labour income shares. The services

FIGURE 3.6: Aggregate labour share approximation

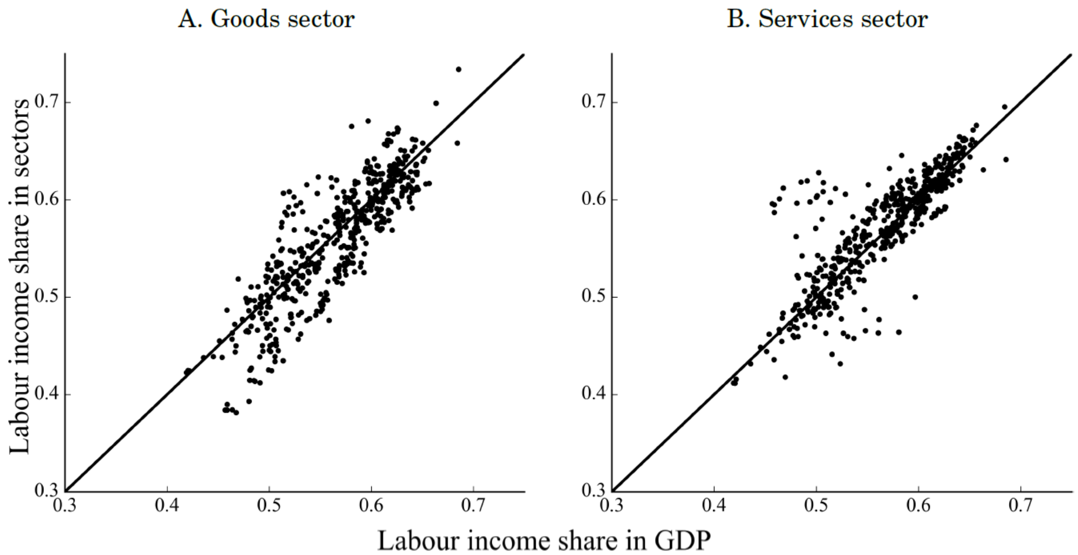

Notes.- aggregate labour income shares (horizontal axes) and sectoral labour income shares (vertical axes) for each year and country. The solid line represents 45 degrees.

sector labour income share is relatively better approximated than the goods sector share, owing to the fact that the expenditure vector places more weight on the relatively larger services sector in the WIOD sample of countries. The aggregate labour income share in each country also provides a better approximation to its sectoral labour income shares than the corresponding sectoral U.S. values.

Therefore, even if no data are available on the input-output structure or value-added labour income shares, it still seems preferable to approximate sectoral labour income shares by the aggregate labour income share within each country, rather than to assume sectoral shares are identical to their corresponding U.S. values across countries.

3.7 Development accounting at the sectoral level

To assess the quantitative importance of the variation in labour income shares across countries, I conduct a development accounting exercise using a framework similar to Herrendorf and Valentinyi (2012). I assume that markets are competitive and production factors are mobile across sectors. However, in contrast to these authors, I do not assume that sectoral labour income shares are identical across countries given the evidence presented above.

Suppose sector $z \in \{G, S\}$ in country i produces commodity output $y_z^{(i)}$ according to a Cobb-Douglas production function:

$$y_z^{(i)} = A_z^{(i)} \left(k_z^{(i)}\right)^{1-\alpha_z^{(i)}} \left(h_z^{(i)}\right)^{\alpha_z^{(i)}} . \tag{3.9}$$

The sectoral quantities of physical capital and human capital are $k_z^{(i)}$ and $h_z^{(i)}$, respectively. Total factor productivity (TFP) of sector z in country i is $A_z^{(i)}$. The parameter $\alpha_z^{(i)} \in (0, 1)$ can vary across sectors and countries.

To implement this accounting exercise, it is necessary to obtain sectoral quantities of y_z, k_z, and h_z for each country. Because the quantities of factor inputs at the expenditure-side are unobserved, I use first-order conditions to obtain these quantities, as in Herrendorf and Valentinyi (2012). Assume the sectoral production function exhibits constant returns to scale and let production factors be mobile across sectors. Let the production function for a given country and for output of sector z be

$$y_z = F_z(k_z, h_z) . \tag{3.10}$$

In this expression, k_z is physical capital and h_z denotes human capital input in sector z. Note that it is not necessary to restrict F to be identical across sectors, and that this method does not require the production function to be of the Cobb-Douglas form. The first-order conditions of a cost minimising representative firm in sector z, combined with the constant returns to scale assumption regarding F, imply

$$s_{k_z} = \frac{p_z(\partial F_z/\partial k_z)k_z}{p_z y_z} = \frac{r k_z}{p_z y_z} , \tag{3.11}$$

where s_{k_z} denotes the capital income share in sector z, p_z denotes the price of sector z's output, and r is the nominal rental rate for capital. It follows that the ratio of capital income across any two sectors z and j is given by

$$\frac{s_{k_z} p_z y_z}{s_{k_j} p_j y_j} = \frac{p_z k_z (\partial F_z/\partial k_z)}{p_j k_j \left(\partial F_j/\partial k_j\right)} = \frac{r k_z}{r k_j} = \frac{k_z}{k_j} , \tag{3.12}$$

due to capital being fully mobile across sectors. Summing over all sectors j, the sectoral share in the aggregate quantity of capital is equal to this sector's share in total capital compensation

$$\frac{k_z}{\sum_j k_j} = \frac{s_{k_z} p_z y_z}{\sum_j s_{k_j} p_j y_j}. \tag{3.13}$$

The sectoral allocation of human capital inputs follows similarly.

The PWT provide data on the stock of physical capital for each country in the WIOD (Feenstra et al., 2015). I construct the measure of aggregate human capital as follows: from the WIOD SEA I obtain the number of total hours worked in each country. Then I multiply this number by an index of human capital, provided in the PWT, which is based on the average years of schooling from Barro and Lee (2013) and the returns to schooling from Psacharopoulos (1994).

To derive quantities of sectoral output, I transform their nominal expenditure values into constant cross-country prices using purchasing power parities. The International Comparison Program (ICP) of the World Bank collects data on prices of a variety of goods and services from almost every country, including all countries in the WIOD, based on national surveys (World Bank, 2008). I use the 2005 version, since this is the most recent version available that lies within the sample period 1995-2009.[6] For a given country, denote nominal expenditure on final use c by e_c, and define the share of expenditure on final use c in total expenditure on sector z as $\lambda_c = e_c / \sum_{c \in z} e_c$. The relevant purchasing power parity in International Dollar for sectoral output z is:

$$PPP_z = \Pi_{c \in z} P_c^{\lambda_c} \,, \tag{3.14}$$

where P_c is the purchasing power parity for final use c. The classification of ICP expenditure categories and the associated sectors is provided in Appendix Table C.2. Finally, I obtain sectoral output quantities by dividing final expenditure on sector z by PPP_z, for every country in the WIOD.

3.7.1 Development accounting results

I compute the TFP of sector z in country i, $A_z^{(i)}$, using equation (3.9) and the quantities as described above. Figure 3.7A-B displays the results for the TFP of the goods and the services sector, using the country and sector-specific labour income shares, and Figure 3.7C-D shows the results using the corresponding U.S. values for sectoral labour income shares. The dashed lines indicate the normalised U.S. productivity level and the log of real GDP per person in the U.S. in International Dollar, which is constant since the U.S. labour income share is identical for both methods.

As expected, sectoral TFP is increasing with the level of development. Comparing panels A and C of Figure 3.7 suggests that productivity differences across countries in the goods sector are slightly smaller when the U.S. labour income shares are used. Similarly, productivity of the services sector seems to increase less in panel D relative

[6] Source: World Bank's ICP 2005, available at http://databank.worldbank.org/...(icp)-2005; accessed 12/04/17.

FIGURE 3.7: Sectoral productivity relative to corresponding U.S. values, all countries 2005: comparison of results using country-specific income shares and U.S. income shares

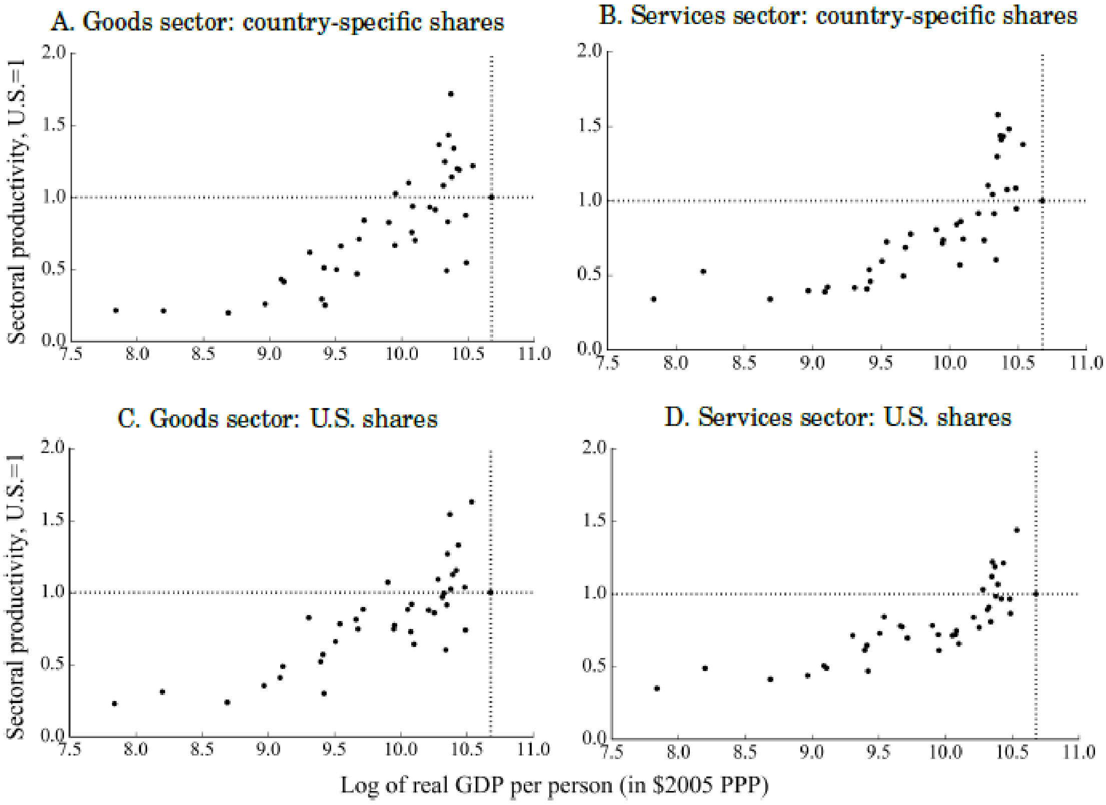

Notes.- vertical axes show real GDP per person in International Dollar in 2005. The horizontal axes show the level of sectoral productivity in 2005, normalised such that U.S.=1. Each dot represents a different country in the WIOD, Luxembourg and is omitted from this figure, but included in the regression.

to panel B of Figure 3.7 with the level of development. To quantify the results of this comparison, I follow Herrendorf and Valentinyi (2012) and report the results of this development accounting exercise in a succinct form: I regress the log of $A_z^{(i)}$ on the log of real GDP per worker in International Dollar of country i, $Y^{(i)}$. Because I obtain $Y^{(i)}$ directly from the PWT, varying the labour income shares does not affect these values.

$$\log\left(A_z^{(i)}\right) = c_z + \gamma_z \log\left(Y^{(i)}\right) + \varepsilon_z^{(i)} , \tag{3.15}$$

where c_z is a sector-specific intercept, γ_z is the elasticity of the TFP of sector z with respect to the real GDP per person, and $\varepsilon_z^{(i)}$ is the residual which captures remaining variation in sectoral productivity not related to aggregate output per worker. A larger value of the estimate $\hat{\gamma}_z$ indicates larger differences of TFP in sector z across countries on average. Regression (3.15) implies:

$$\frac{A_z^{(i)}}{A_z^{USA}} = \left(\frac{Y^{(i)}}{Y^{USA}}\right)^{\gamma_z} \exp\left(\frac{\varepsilon_z^{(i)}}{\varepsilon_z^{USA}}\right). \tag{3.16}$$

Therefore, if $\hat{\gamma}_z$ equals one, then the difference in aggregate labour productivity between the U.S. and other countries is on average as large as the difference between sectoral TFP of those countries in sector z. If $\hat{\gamma}_z$ equals zero, then there is no systematic relationship between aggregate labour productivity and the TFP of sector z across countries. Table 3.4 displays the estimates of $\hat{\gamma}_z$ for the goods and the services sector of all countries in the WIOD.

TABLE 3.4: Parameter estimates of $\hat{\gamma}_z$, all countries 2005

	Country and sector-specific (1)	U.S. (2)	Fixed = 0.66 (3)
1. Aggregate economy	0.66***	0.55***	0.64***
	(0.05)	(0.04)	(0.04)
2. Goods sector	0.77***	0.66***	0.76***
	(0.06)	(0.05)	(0.05)
3. Services sector	0.55***	0.46***	0.57***
	(0.05)	(0.04)	(0.04)

Notes.- column (1) shows sectoral TFP using sector- and country-specific labour income shares. Column (2) assumes sectoral labour income shares equal their corresponding U.S. values across countries. Column (3) assumes a fixed value equal to 0.66 for sectoral labour income shares. Newey-West standard error estimates in parenthesis.
*** p<0.01, two-sided test.

Column (1) of Table 3.4 shows the estimates of $\hat{\gamma}_z$ when sectoral productivity levels are computed using country and sector-specific labour income shares. The estimate of 0.66 in the first row shows that the aggregate TFP difference increases on average by 0.66 per cent when the difference in aggregate output per worker increases by one per cent. The sectoral TFP differences across countries are on average larger in the goods sector than in the aggregate economy. In contrast, systematic TFP differences in the services sector are smaller than aggregate TFP differences. These results support the findings of Herrendorf and Valentinyi (2012), who also document larger TFP disparities in the goods sector than in the services sector.

Table 3.4 column (2) displays the estimates when I impose that sectoral labour income shares are equal to the corresponding U.S. value across countries. The qualitative results are unchanged: productivity differences are larger in the goods sector than in the services sector across countries. However, the magnitudes of all coefficients are reduced by almost 20 per cent, which implies that the association between aggregate labour productivity and sectoral TFP is on average weaker. Relatively poor countries have on average lower ratios of physical to human capital. Therefore, decreasing the exponent of human capital in the production function (3.9) decreases the part of the observed output difference which can be explained by physical capital intensity. This means that the accounting residual, TFP, has to increase across countries.

The third column of Table 3.4 shows the estimates when the sectoral labour income shares are equal to two-thirds, a common value for the labour income share found in the literature. Not surprisingly, the qualitative results from columns (1) and (2) are unchanged, and the estimates are resembling those from the first column, because a relatively high labour income share increases the dispersion on the left hand side of regression (3.16). Coefficient estimates in this column are relatively better approximations of the benchmark estimates in column (1), suggesting that imposing U.S. labour income shares across countries provides a poor approximation.

Less developed countries on average have lower stocks of physical and human capital than more developed countries. Therefore, using the relatively high U.S. values for labour income shares across countries overestimates the contribution of human capital in the production process. Using the country and sector-specific labour income shares thus increases the productivity differences across countries, since the input factors account for less of the observed differences in aggregate output.

3.8 Conclusion

The expenditure-side labour income shares of the goods and services sectors vary substantially across countries and increase significantly with the level of development. This result is robust to different adjustment methods to account for the labour income of the self-employed. Therefore the usual assumption that sector-specific labour income shares are identical across countries to their corresponding U.S. values is not justified.

I use a decomposition to show that the observed differences in sectoral labour income shares across countries are mainly driven by differences in labour income at the production-side of the economy. The input-output structure and expenditure weights do not contribute substantially to the variation across countries in the labour income shares of the goods sector. Moreover, expenditure weights and input-output linkages even decrease the variation in the labour income shares of the services sector across countries.

The accounting method used in this paper to compute the labour income shares at the sectoral level require relatively detailed data on the structure of an economy, and thus it is not applicable to countries who lack statistical capacity. Therefore, I show that a less data-intensive approximation of the sectoral labour income shares by the value-added weighted production-side labour income shares provides a good fit. Moreover, even the aggregate labour income share provides a reasonably close approximation, superior to the fit achieved by imposing U.S. labour income shares across countries.

The findings here also have implications for development accounting exercises: cross-country productivity differences are larger than previously thought. Computing sectoral TFPs using labour income shares which increase in the level of development increases the accounting residual, because less developed countries have a lower capital intensity. I find that the goods sector exhibits relatively larger productivity differences across countries than the services sector.

Conclusion

The first essay provided new estimates of the flexibility of real wages during the UK's Great Recession. The novelty of my findings is that I was able to control for cyclical job-switching, which might have obscured I found that job-stayer real wages respond by as much as 2.6 per cent for every one percentage point rise in the unemployment rate. Hiring wages are at least as responsive to the business cycle as the wages of job stayers. I also found that the hours worked by job stayers did not respond to the Great Recession. Conversely, the hours of new hires among the same firms responded significantly, decreasing by 1.5 per cent for every one percentage point rise in the unemployment rate, mostly through firms switching between full- and part-time workers.

Cohorts hired during the Great Recession were not only paid lower wages initially, but were also locked into low-wage growth paths. This significantly reduced the present value of labour costs from the firm's perspective for hires made during this time, *ceteris paribus*. In this respect, it seems that firms' hiring wages were even more flexible than the results for the initial real wages of new hires indicate. Moreover, when combined with the shift from full- to part-time hiring, firms were able to significantly reduce their labour costs per new employee. While these large and significant wage and hours responses seem very likely to account for the UK economy's unusual experience of the Great Recession, the puzzle still remains as to why firms were able to adjust wages so freely, and why workers were so willing to accept these changes. This is potentially a fruitful area for future research.

In the second essay, I showed that relaxing the assumption of rational expectations in the textbook search and matching model of the labour market does not affect the equilibrium level of market tightness, wages, and unemployment. In particular, assuming that agents use linear forecasts does not only result in local convergence but also extends trivially to global convergence for all reasonable parameter values usually found in the literature. So, the potentially very restrictive rational expectation assumption in this class of models is nonetheless reasonable with respect to equilibrium outcomes. I used recent UK data to parametrise the model, and showed that implied convergence to the equilibrium can be very slow, however. Therefore, the rational expectation model of unemployment fluctuations could in fact be a poor approximation to the dynamics of an economy in which agents learn as econometricians, especially in the presence of frequent structural or permanent shocks.

In the final essay, I found that the expenditure-side labour income shares of the goods and services sectors vary substantially across countries and increase significantly with the level of development. This result is robust to different adjustment methods to account for the labour income of the self-employed. Therefore the usual assumption that sector-specific labour income shares are identical across countries to their corresponding U.S. values is not justified for the countries studied, which together account for nearly 85 per cent of the global output.

These findings also have implications for development accounting exercises: cross-country productivity differences are larger than previously thought. Computing sectoral TFPs using the here computed labour income shares, which increase in the level of development, increases the accounting residual, because less developed countries have a lower capital intensity. I find that the goods sector exhibits relatively larger productivity differences across countries than the services sector, providing evidence against the claim of some studies that the services sector is relatively less productive.

Appendix A

Real wages and hours in the Great Recession: Evidence from firms and their entry-level jobs

Appendix A.1 Further description of the data and sample

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

We focus on methodological details through the period 1998-2016. Throughout this period, the ASHE should be a true random sample of all employees in employment, irrespective of employment status, occupation, size of employer etc. Given the legal obligation of employers to respond using payrolls, it has a high response rate and is believed to be accurate. There is no cumulative attrition from the panel, as any individual not included in the ASHE in any year, for whatever reason, remains in the sampling frame the following year. Conditional on a hundred per cent response, the ASHE is a true one per cent random sample of employees: all with a National Insurance Number which has a numerical part ending in 14. However there are two major sources of under-sampling, both occurring if individuals do not have a current tax record. This could happen for some individuals who have very recently moved job, or for those who earn very little (mostly part-time) and are not paying income tax or National Insurance in the period when their employers are looked up. From

2004 the ASHE aimed to sample some of those employees under-represented. It added supplementary responses for those without a PAYE reference, and also attempted to represent employees whose jobs changed between the determination of the sampling frame in January and the reference period in April. Since the ONS states that the biases that these amendments were introduced to address were actually small, we do not believe they could affect our results substantially. The ASHE also introduced some imputations, using similar matched 'donor' observations where responses were, for example, missing an entry of basic hours but had recorded pay. These imputations were added for weighting purposes, but throughout our analysis we ignore the weights in the ASHE, since they are designed to make the aggregate results population representative in terms of worker observables, and are not firm-level.

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

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

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

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

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

2003 codes is decided using whichever of the two bore a greater share of economic output.

Appendix A.2 Further robustness checks

Table A.1 presents some further robustness checks of the main empirical results

TABLE A.1: Estimated semi-elasticity of real wages and hours with respect to the unemployment rate, 1998-2016: more robustness checks

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

Notes.- second-step regression results of estimated period effects on unemployment rate, $\hat{\gamma}$. The first row is identical to Table 1.3, included here for comparison. The second row uses mean wages in jobs as the dependent variable in the first step. The third row excludes all time-varying controls from the first step. The forth row includes public sector firms in the analysis. The fifth row uses the Retail Price Index, instead of the Consumer Price Index, to deflate wages. The sixth row uses less restricted values of the first-step dependent variables: wages include shift-work, incentive payments, overtime, and all other payments; hours refer to basic and paid overtime.

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

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

presented in Table 1.3. The first row repeats our baseline/main result for convenience. The robustness discussed here is with regards to the specification of the first step of the regression model: (1.1) & (1.3). The specification of the second step is unchanged compared with the baseline. The second row describes the estimated semi-elasticity of real wages and hours with respect to the unemployment rate when typical job-level measures are employee sample means, rather than median values. Qualitatively the results are unchanged: wages for hires and job stayers exhibit a sizeable and significant cyclical response, as do hiring hours, though the difference in wage response between hires and job stayers is larger. We prefer the median as a measure of the typical wage because it is less sensitive to changes in the extent of sampling error within jobs over time, given our specific sample selection criteria for jobs. The third row removes all controls for time-varying job characteristics from the first step. In doing so we would expect to underestimate the cyclical response of wages because of a procyclical composition bias along some observable characteristics. However the results here show that those observables that we do control for at the job level, namely gender, union coverage, age and firm size, are collectively not important in this regard. The fourth row includes jobs from the public sector. The main findings are qualitatively unchanged. The hiring hours in public sector entry-level jobs were somewhat more responsive to the Great Recession than in the private sector, potentially reflecting the squeeze on labour costs imposed by fiscal austerity. The fifth row simply illustrates the difference in results when we use an alternative price deflator. The RPI notably includes the cost of housing, including mortgage interest payments, whereas the CPI does not. Interest rates were cut during the Great Recession, and so the RPI is itself more cyclical than the CPI. Hence the measured real RPI-wage cyclicality is smaller, though still significant. We prefer the CPI because it is more internationally comparable and is the basis of the Bank of England's inflation target. The sixth row includes other work-related payments in earnings and the derived hourly wage rate, such as incentive or overtime pay. It similarly includes overtime and shift work in hours worked. There are reasonable arguments why including these other payments could lead to both increased or decreased wage responsiveness. Here their inclusion has no significant effect, except for job stayers' hours, suggesting that working hours for employees with overtime and shift work are more responsive than standard working hours.

TABLE A.2: Estimated semi-elasticity of real wages and hours with respect to the unemployment rate, 1998-2016: sample selection robustness - varying the minimum number of employees per job-year required for inclusion in the CH-firms sample

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

Notes.- second-step regression results of estimated period effects on unemployment rate, $\hat{\gamma}$. Each row gives results varying the minimum number of employees per job-year required for selection into the analysis sample. "3 (baseline)" is identical to Table 1.3, included here for comparison.

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

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

TABLE A.3: Estimated semi-elasticity of the difference in $\hat{\beta}_t$ series between new hires and job stayers with respect to the unemployment rate, 1998-2016

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

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

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

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

Appendix A.3 Additional tables

TABLE A.4: Distribution of new hires over industries, all years 1998-2016

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

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

TABLE A.5: Distribution of new hires over occupations, all years 1998-2016

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

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

TABLE A.6: Estimated period-fixed effects for real hourly wages ($\hat{\beta}_t$ from first-step regressions)

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

Notes.- time series of period-fixed effects for different subsamples of the ASHE. Estimated using (1.1) and (1.3). Normalised to zero in 2003. (1) Entry-level new hires, (2) job stayers in CH-firms (3) ASHE new hires, (4) ASHE job stayers.

TABLE A.7: Estimated period-fixed effects for basic weekly hours worked ($\hat{\beta}_t$ from first-step regressions)

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

Notes.- see Table A.6

TABLE A.8: Time series of price deflators and business cycle indicators

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

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

TABLE A.9: Estimated hiring-year-tenure-fixed effects for real wages and hours relative to their hiring levels: workers who stay in entry-level jobs

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

Notes.- see Section 1.5, Figure 1.6. Ordinary least squares standard errors in parentheses.

Appendix A.4 Additional figures

FIGURE A.1: Characteristics of employees in the consistent-hiring-firms sample and whole ASHE: comparison of new hires in entry-level jobs vs. job stayers, 1998-2016

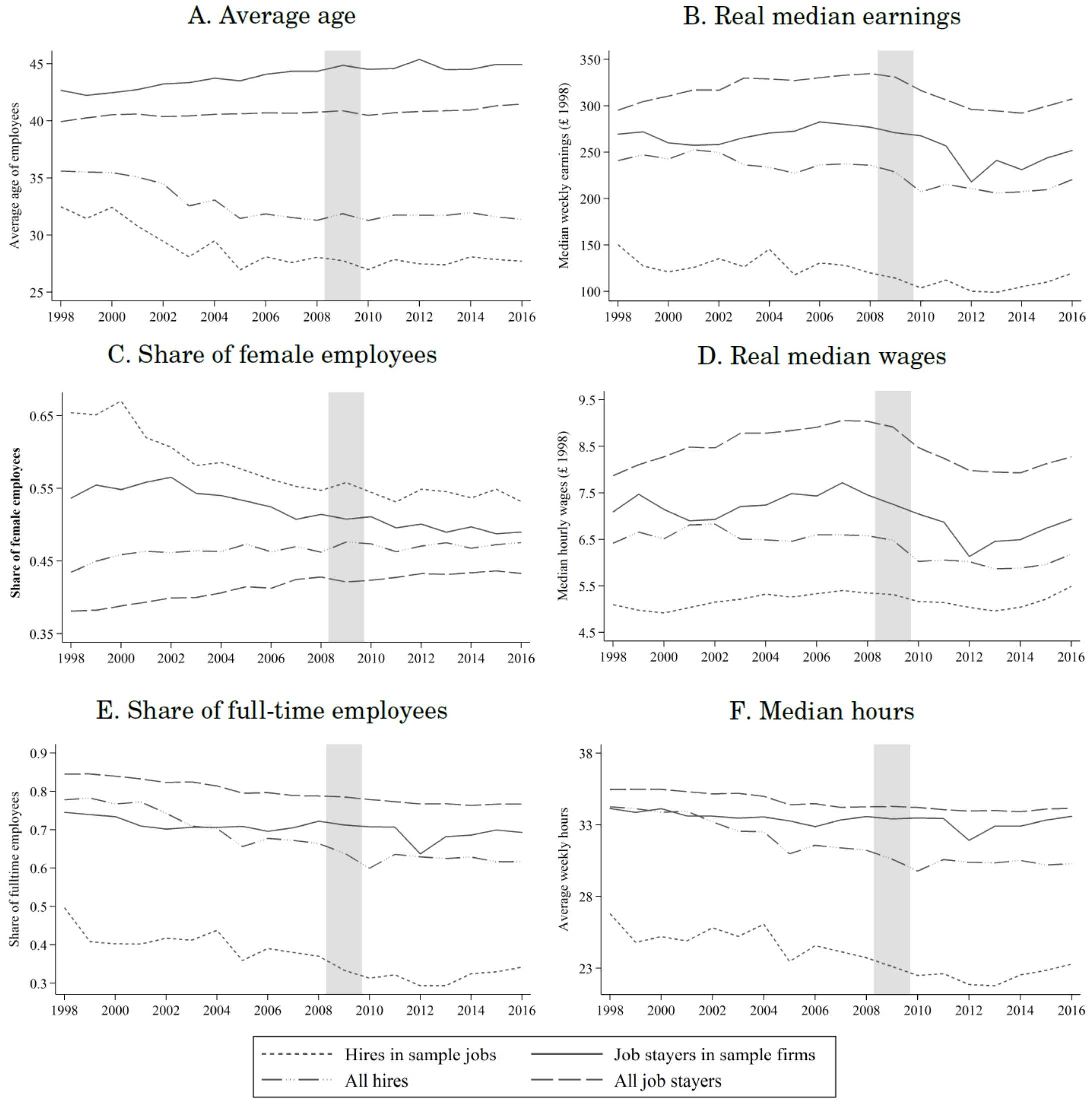

Notes.- shaded area marks official UK recession dates. “Hires in sample firms" refers to employees in entry-level jobs with less than twelve months tenure. “Job stayers in sample firms" are for jobs and employees who have more than 12 months tenure in the same job, and only for firms which are represented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

FIGURE A.2: Comparison of business cycle indicators, 1998-2016

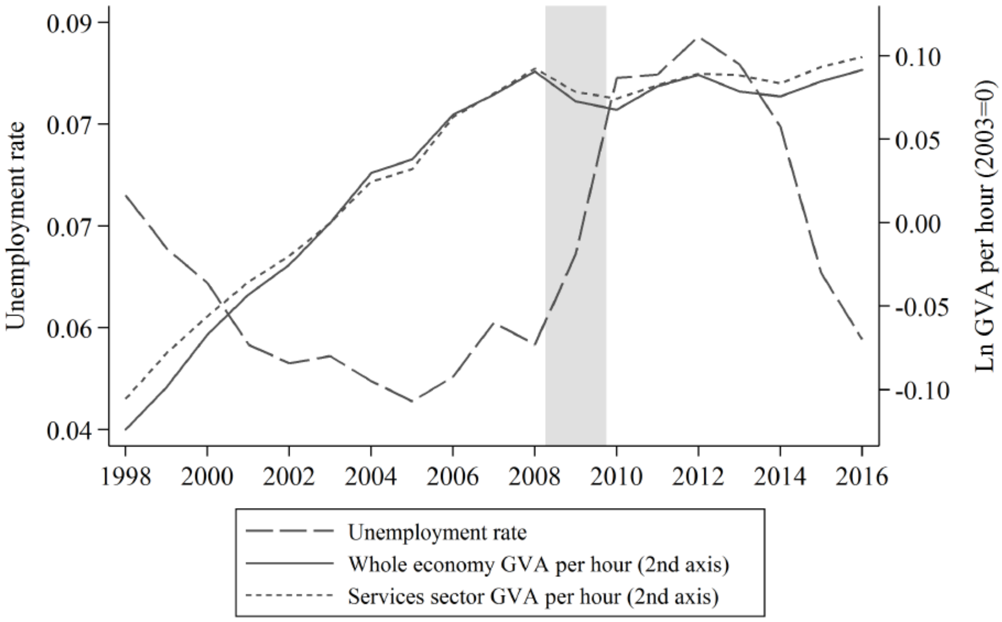

Notes.- see Section 1.2 and Table 1.5 for sources. Shaded area marks official UK recession dates.

FIGURE A.3: Estimated period-fixed effects for real wages, 1998-2016: comparison of entry-level jobs, all new hires, job stayers in the CH-firms sample, and all job-stayers

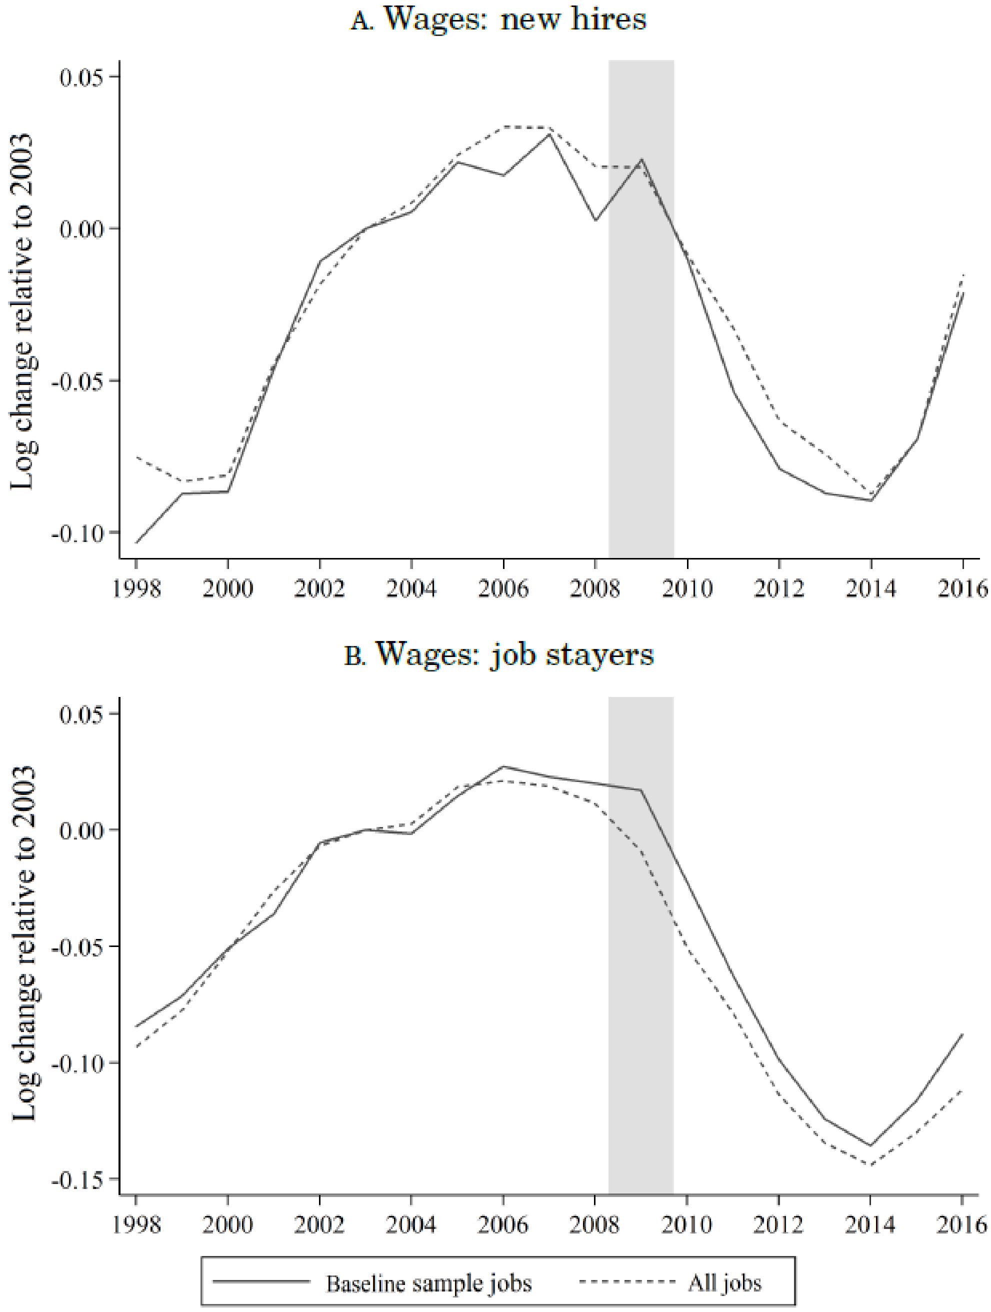

Notes.- see Figure A.2 and Section 1.2 for further details of sample construction. “All” here refers to all firms and jobs represented in the ASHE. Shaded area marks official UK recession dates.

FIGURE A.4: Estimated period-fixed effects for hours worked, 1998-2016: comparison of entry-level jobs, all new hires, job stayers in the CH-firms sample, and all job-stayers

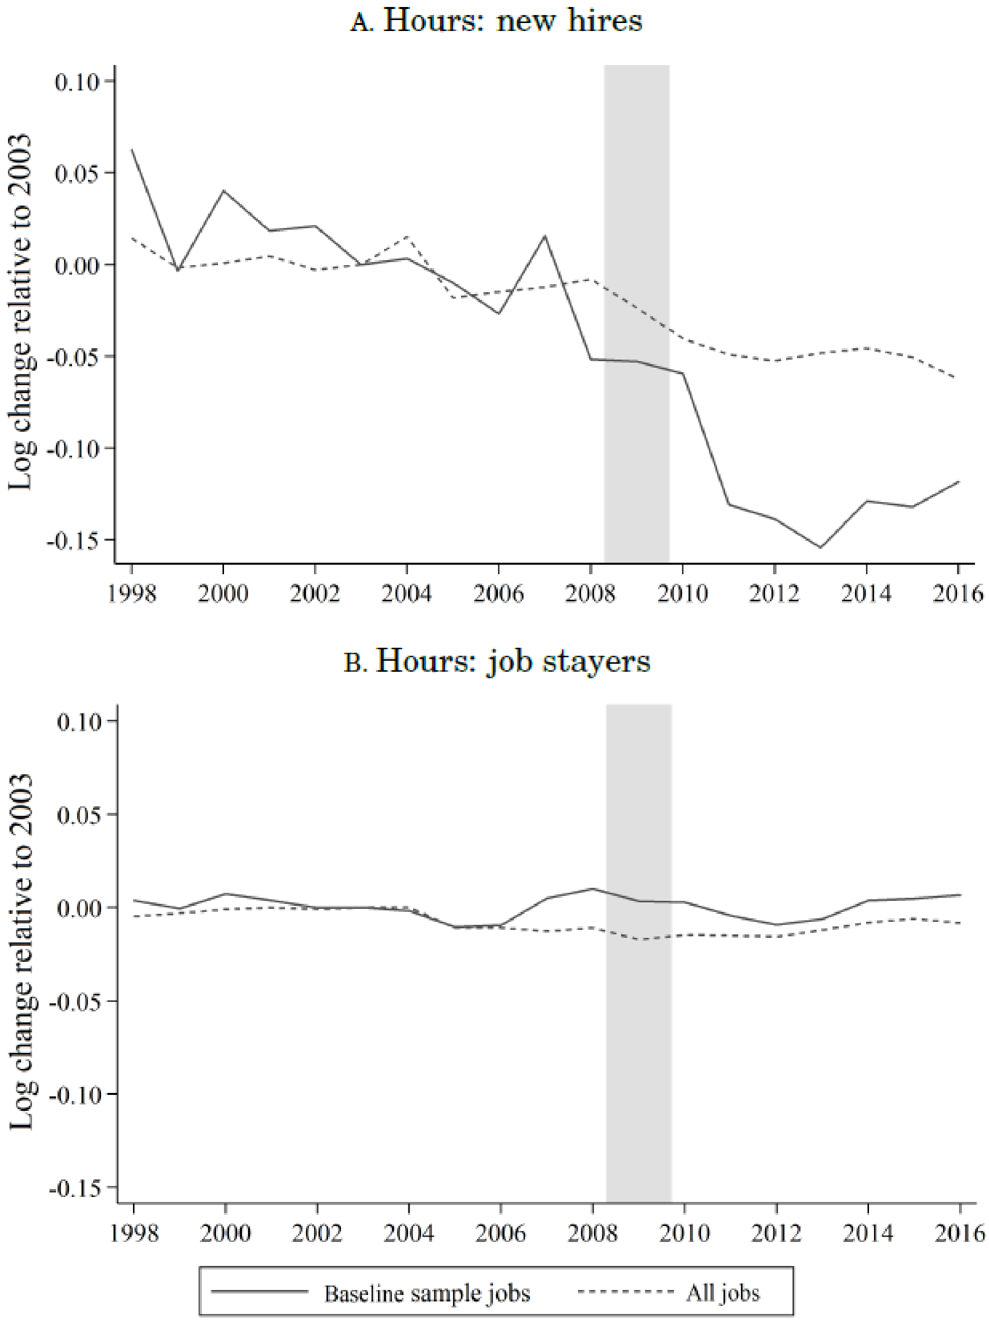

Notes.- see Figure A.2 and Section 1.2 for further details of sample construction. “All” here refers to all firms and jobs represented in the ASHE. Shaded area marks official UK recession dates.

FIGURE A.5: Estimated period-fixed effects for real wages including linear trend, 1998-2016: comparison of new hires in entry-level jobs vs. job stayers

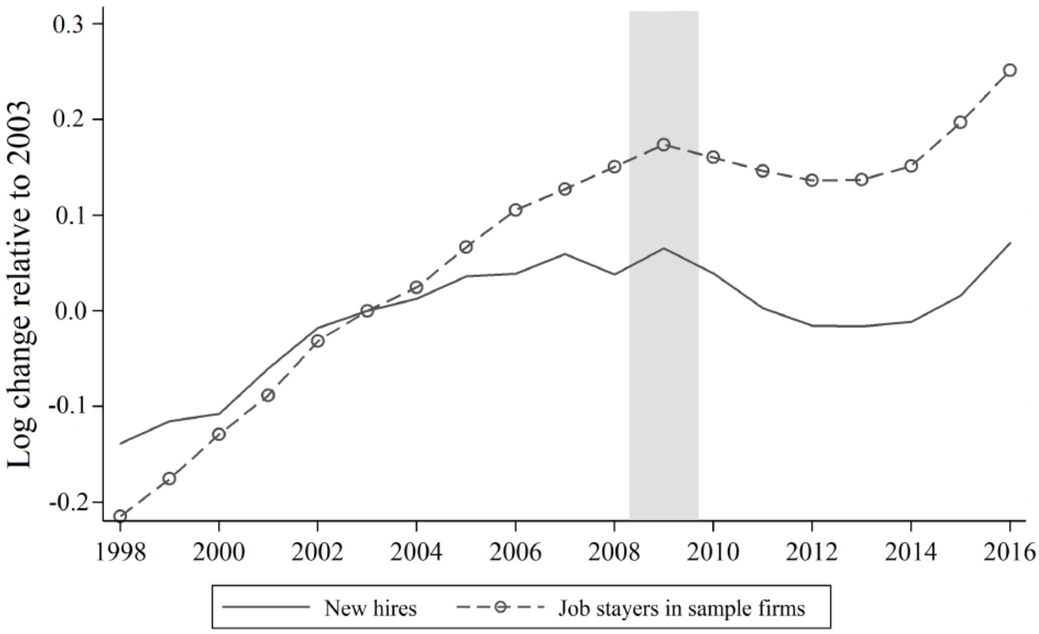

Notes.- see Figure 1.2, except here the series are not adjusted for linear trends. 2003 normalised to zero.

Appendix A.5 Nominal wage changes of job stayers in CH-firms

FIGURE A.6: Distribution of year-to-year changes in log nominal hourly wages for job stayers in CH-firms, 2005-06 and 2009-10

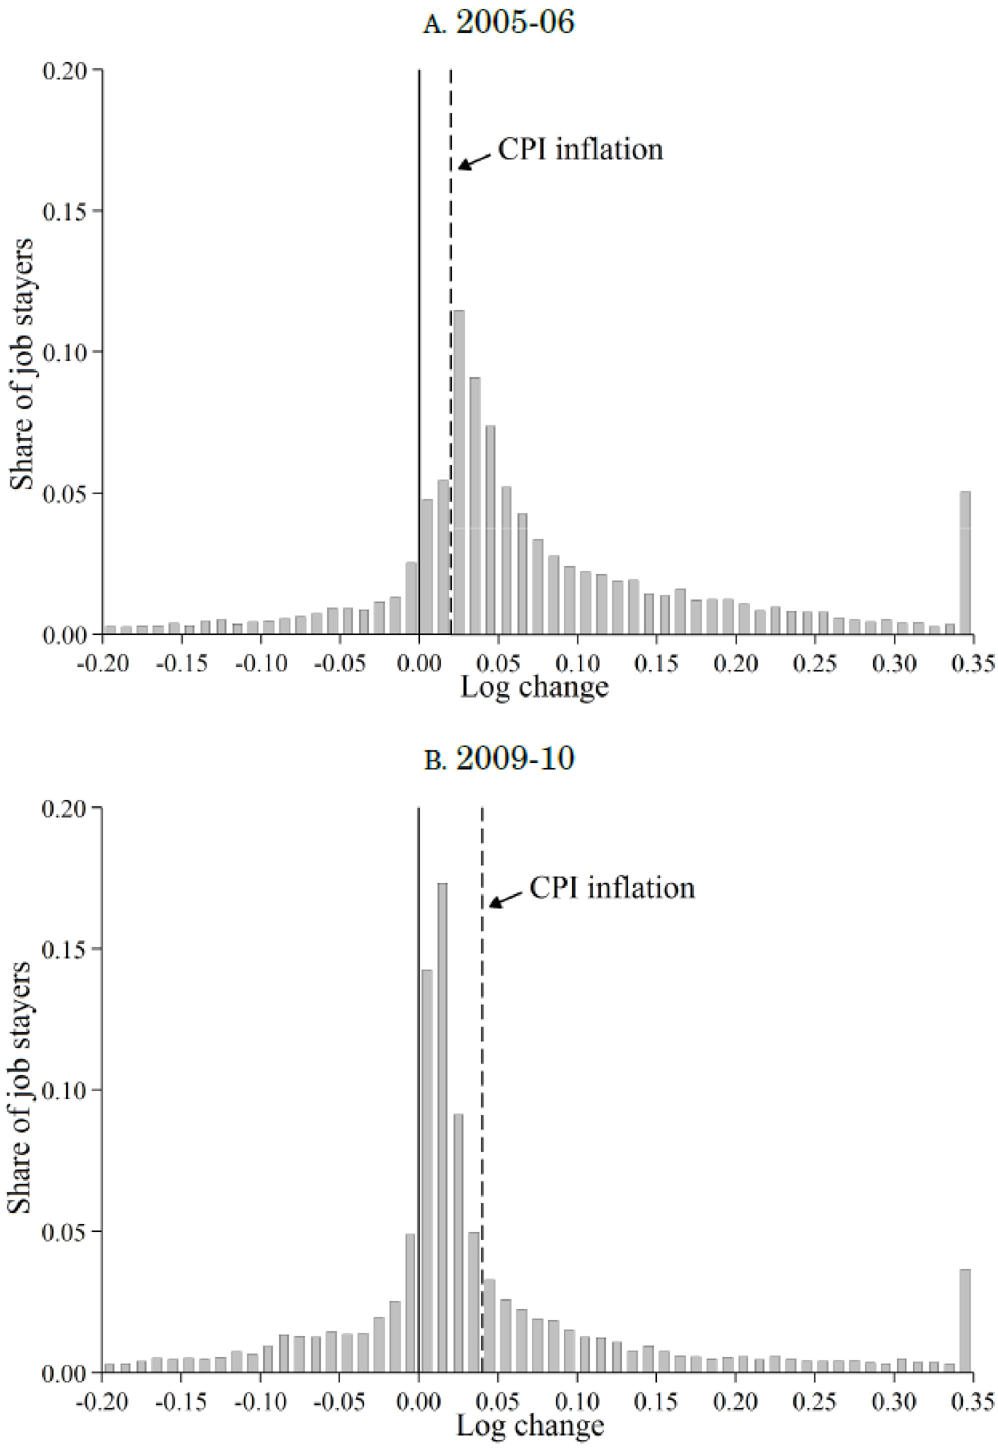

Notes.- solid line marks zero, dashed line marks the log change in the Consumer Price Index. Bars show half-open intervals, excluding the upper limit. Ages 16-64, private sector only.

TABLE A.10: Percentage of job stayers with year-to-year change in log nominal hourly wages in given category, 1997-2016

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

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

Nominal wage changes of job stayers in the UK have been analysed previously by Nickell and Quintini (2003) and most recently by Elsby et al. (2016). We briefly summarise results for year-to-year changes in the log nominal wages of job stayers in our baseline sample of firms. Figure A.6 shows the distributions of log changes for job stayers between 2005-06 and 2009-10. These two periods are representative of periods with relatively low (2005-06) and relatively high (2009-10) shares of job stayers with nominal wage cuts, see Table A.10 which displays summary statistics for all years in our sample. The dashed line marks the inflation rate in the histograms. Bars in the histograms exclude upper limits, so log wage changes of exactly zero are included in the bin to the right of the solid line.

The spike at zero is relatively small during normal times, ranging from 0.5 per cent to 2.1 per cent in the period before the Great Recession as Table A.10 shows. The distribution in panel A also suggests that most wages increase with the rate of inflation during normal times, thus keeping the real wage constant. Nevertheless, even during this period a notable share of job stayers, around 20 per cent, experienced nominal wage cuts. This share increased during the recession to around 25 per cent on average. Similarly, the share of job stayers with exactly zero nominal wage growth peaked at

5.7 per cent between 2009-10. In particular, Figure A.6B displays a relatively large share of nominal wage changes between zero and two per cent for job stayers in CH-firms between 2009-10. However, the large increase in the percentage of job stayers which experienced negative changes in log *real* wages, as shown in the last column of Table A.10, was mainly caused by the rise in inflation.

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

Appendix A.6 Description of the kernel re-weighting method

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

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

$$g(w|t=04;m_{04}) = \int_{\Omega_x} f^{04}(w|x;m_{04})h(x|t=04)dx \ , \tag{A.1}$$

and

$$g(w|t=13;m_{13}) = \int_{\Omega_x} f^{13}(w|x;m_{13})h(x|t=13)dx \ , \tag{A.2}$$

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

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

$$g(w|t=13;m_{04}) = \int_{\Omega_x} f^{13}(w|x;m_{04})h(x|t=13)dx \ , \tag{A.3}$$

where we know $h(x|t=13)$, but the density of prices $f^{13}(w|x;m_{04})$, consisting of the real wage schedule of 2013 and the real minimum wage of 2004, is unobserved. We can partition the density in (A.3) into the part of real wages below the real value of the minimum wage in 2013 and the part of real wages above this threshold:

$$\begin{aligned} g(w|t=13;m_{04}) = &\int_{\Omega_x} [1-I(w \le m_{13})]\, f^{13}(w|x;m_{04})h(x|t=13)dx \\ &+ \int_{\Omega_x} I(w \le m_{13}) f^{13}(w|x;m_{04})h(x|t=13)dx \ , \end{aligned} \tag{A.4}$$

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

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

$$[1 - I(w \leq m_H)]\, f^i(w|x; m_H) = [1 - I(w \leq m_H)]\, f^i(w|x; m_L)\ . \tag{A.5}$$

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

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

$$I(w \leq m_H) f^H(w|x; m_L) = I(w \leq m_H) \psi_w f^L(w|x; m_L)\ , \tag{A.6}$$

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

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

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

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

We illustrate the method in Figure A.7, which shows estimated densities of real hiring wages (we pool data over multiple periods in Figures A.7 and A.8 for data confidentiality reasons). The kernel density estimate in the top panel uses pooled data from 2002-04, a period where the NMW was relatively less binding. The solid vertical line shows the real value of the NMW at its 2013 level. Panel B shows the corresponding density for 2012-14, where the real NMW remained nearly constant around its 2013 value. The counterfactual density displayed in panel C is a simple combination of the part of the density to the left of the solid line in panel A and to the right of the solid line in Panel B, scaled to integrate to one.

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

FIGURE A.7: Illustration of the re-weighting procedure for log real hourly hiring wages

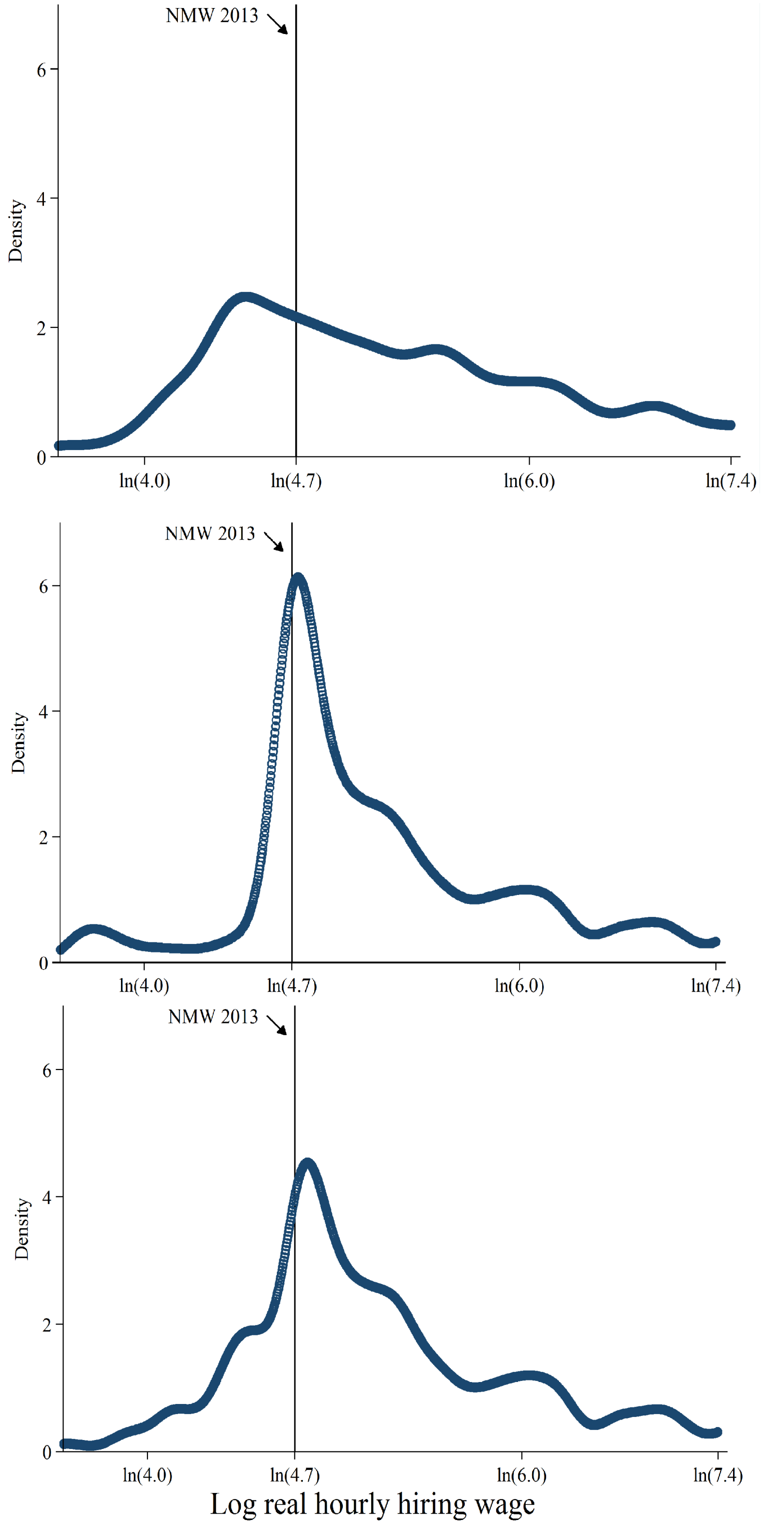

Notes.- densities estimated using Gaussian kernel and bandwidth of 0.03 (A) and 0.02 (B-C). Monetary values deflated to 1998 values using the CPI. Solid lines show the real value of the adult rate minimum wage in 2013.

Figure A.8D plots the estimated (connected circles) and counterfactual (solid line) density for 2012-14 together. Most of this mass originates from jobs which are observed to hire slightly above the NMW, a result of the smoothing by the kernel estimator. For hiring wages which exceed the NMW substantially, the estimated and counterfactual density are, as expected, indistinguishable. Figure A.8E displays the difference between the estimated and counterfactual wage densities shown in Panel D of this figure. The difference is negative for values around the value of the NMW in 2013 and positive for log hiring wages between $\ln(4)$ and $\ln(4.6)$.

Assumptions 1-3 allow us to write

$$\begin{aligned} g(w|t=13;m_{04}) = {} & [1-I(w \le m_{13})]\, f^{13}(w|x;m_{04})dx \\ & I(w \le m_{13})\,\psi_w f^{04}(w|x;m_{04})dx \ , \end{aligned} \tag{A.7}$$

with

$$\psi_w = \frac{\Pr(w \le m_{13}|x, f = f^{13})}{\Pr(w \le m_{13}|x, f = f^{04})} \ , \tag{A.8}$$

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

$$\begin{aligned} g(w|t=13;m_{04}) = {} & \int_{\Omega_x} [1-I(w \le m_{13})]\, f^{13}(w|x;m_{04})h(x|t=13)dx \\ & + \int_{\Omega_x} I(w \le m_{13})\,\psi_w f^{04}(w|x;m_{04})h(x|t=13)dx \ . \end{aligned} \tag{A.9}$$

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

$$\begin{aligned} g(w|t=13;m_{04}) = {} & \int_{\Omega_x} [1-I(w \le m_{13})]\, f^{13}(w|x;m_{04})h(x|t=13)dx \\ & + \int_{\Omega_x} I(w \le m_{13})\,\psi_w f^{04}(w|x;m_{04})\theta h(x|t=04)dx \ , \end{aligned} \tag{A.10}$$

where the re-weighting function is

$$\theta = \frac{h(x|t=13)}{h(x|t=13)} = \frac{\Pr(t=13|x)}{\Pr(t=04|x)}\frac{\Pr(t=04)}{\Pr(t=13)} \ . \tag{A.11}$$

FIGURE A.8: Estimated and counterfactual densities of real log hiring wages at the job-level, 2012-14

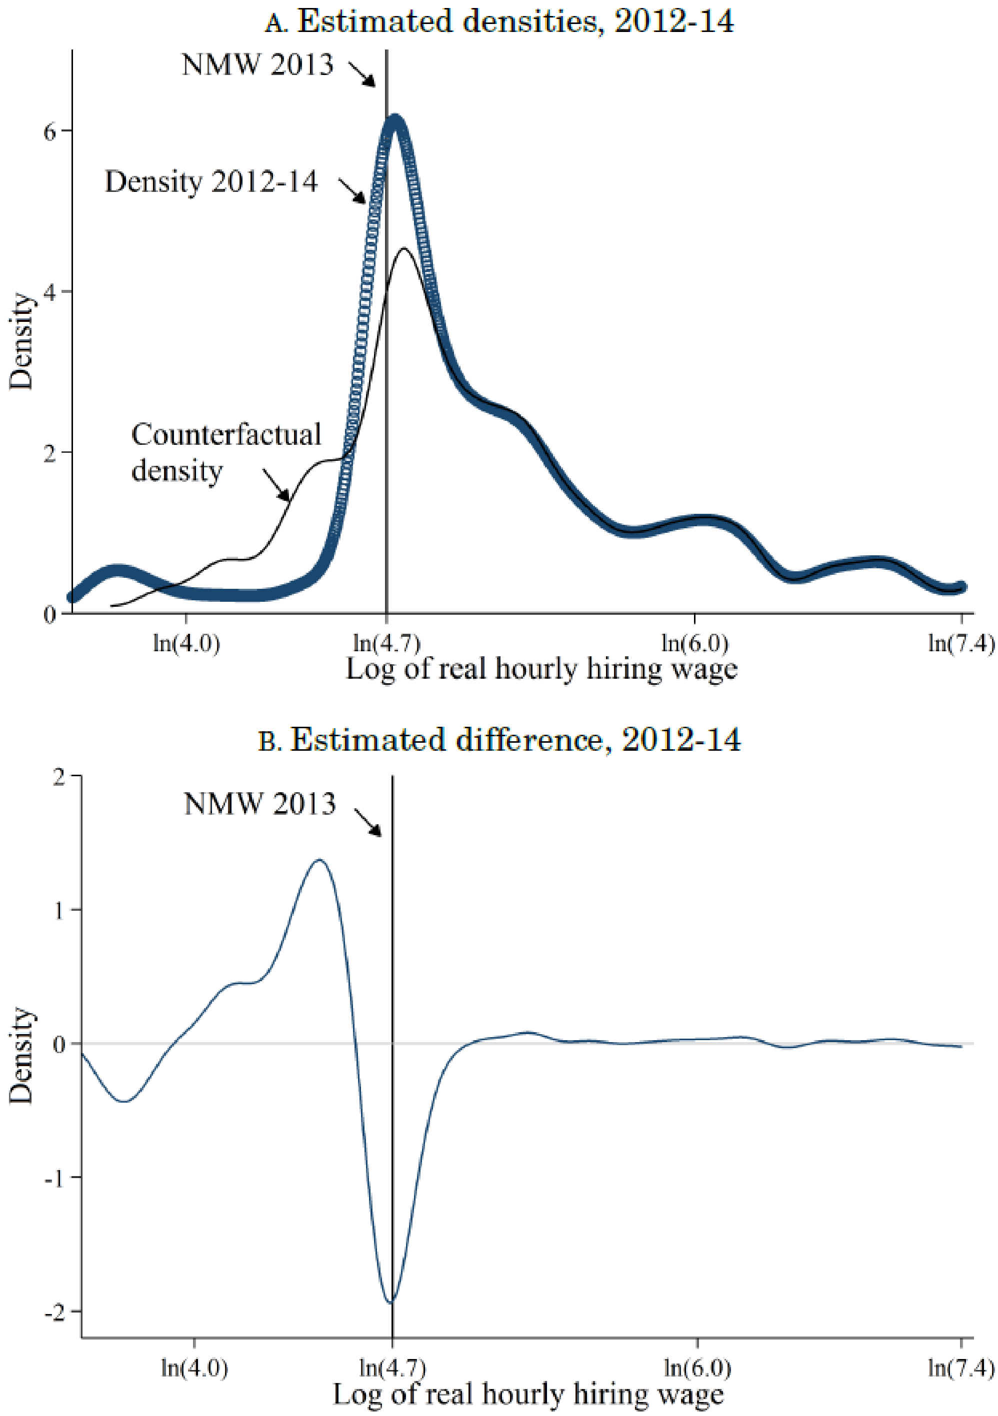

Notes.- densities estimated using Gaussian kernel and bandwidth of 0.02. Monetary values deflated to 1998 values using the CPI. Solid vertical line shows the real value of the adult rate minimum wage.

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

$$\psi = \theta \cdot \psi_w = \frac{\Pr(t = 13|x, w \leq m_{13})}{\Pr(t = 04|x, w \leq m_{13})} \frac{\Pr(t = 04)}{\Pr(t = 13)} . \tag{A.12}$$

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

$$\Pr(t = 13|x, w \leq m_{13}) = \Lambda(C(x)) \ , \tag{A.13}$$

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

Appendix B

Unemployment and econometric learning

Appendix B.1 Methodology

B.1.1 Linearisation

We take a first order Taylor approximation around the deterministic steady-state values of θ and y, $\bar{\theta}$ and $\bar{y} = 1$ respectively, approximating the right and the left hand side of equation (2.20) which is stated here again for convenience,

$$\frac{c}{\delta q(\theta_t)} = \left[(1-\beta)(y_{t+1} - b) + \frac{(1-\lambda)c}{q(\theta_{t+1})} - \theta_{t+1}\beta c\right]^e .$$

This results in

$$\begin{aligned} \frac{c}{\delta q(\bar{\theta})} - \frac{cq'(\bar{\theta})}{\delta[q(\bar{\theta})]^2}(\theta_t - \bar{\theta}) =& (1-\beta)(\bar{y} - b) + (1-\beta)(y^e_{t+1} - \bar{y}) \\ & - \bar{\theta}\beta c - \beta c(\theta^e_{t+1} - \bar{\theta}) + \frac{(1-\lambda)c}{q(\bar{\theta})} \\ & - \frac{(1-\lambda)cq'(\bar{\theta})}{[q(\bar{\theta})]^2}(\theta^e_{t+1} - \bar{\theta}). \end{aligned} \tag{B.1}$$

By noting that

$$\frac{c}{\delta q(\bar{\theta})} = (1-\beta)(\bar{y} - b) + \frac{(1-\lambda)c}{q(\bar{\theta})} - \bar{\theta}\beta c$$

must hold in equilibrium according to (2.20), this steady-state condition can be subtracted from both sides of the approximated equation. Then solving explicitly for

θ_t and defining the functional form $q(\bar{\theta}) = \mu\bar{\theta}^{-\alpha}$, (B.1) becomes

$$\begin{aligned}\theta_t = &\left\{1 + \frac{\beta\delta\mu^2(\bar{\theta})^{-2\alpha}}{\alpha\mu(\bar{\theta})^{-\alpha-1}} - (1-\lambda)\delta\right\}\bar{\theta} \\ &- \frac{(1-\beta)\delta\mu^2(\bar{\theta})^{-2\alpha}}{c\alpha\mu(\bar{\theta})^{-\alpha-1}}\bar{y} \\ &+ \frac{(1-\beta)\delta\mu^2(\bar{\theta})^{-2\alpha}}{c\alpha\mu(\bar{\theta})^{-\alpha-1}}y^e_{t+1} \\ &+ \left\{-\frac{\beta\delta\mu^2(\bar{\theta})^{-2\alpha}}{\alpha\mu(\bar{\theta})^{-\alpha-1}} + (1-\lambda)\delta\right\}\theta^e_{t+1},\end{aligned} \tag{B.2}$$

which can be simplified to the form given in the text, (2.21), with coefficients

$$\begin{aligned}\psi_0 &= [1-\psi_2]\,\bar{\theta} - \psi_1\bar{y}, \\ \psi_1 &= \frac{(1-\beta)\delta\bar{\theta}q(\bar{\theta})}{c\alpha}, \\ \psi_2 &= \delta\left[(1-\lambda) - \frac{\beta\bar{\theta}q(\bar{\theta})}{\alpha}\right],\end{aligned}$$

and with the steady-state value for labour market tightness the solution to

$$(1-\beta)(\bar{y}-b) - \frac{c(\frac{1-\delta}{\delta}+\lambda)}{q(\bar{\theta})} - \beta c\bar{\theta} = 0.$$

B.1.2 Determinacy of the REE

The operator E_t denotes mathematical expectations formed at period t. The linearised dynamics of output (2.22) can be substituted into (2.21) by noting under RE that $E_t y_{t+1} = (1-\rho) + \rho y_t$;

$$\theta_t = \tilde{\psi}_0 + \tilde{\psi}_1 y_{t-1} + \tilde{\psi}_2 E^*_t\theta_{t+1} + \tilde{\psi}_1\rho^{-1}\varepsilon_t, \tag{B.3}$$

with

$$\begin{aligned}\tilde{\psi}_0 &= \psi_0 + \psi_1(1-\rho)(1+\rho), \\ \tilde{\psi}_1 &= \psi_1\rho^2, \\ \tilde{\psi}_2 &= \psi_2.\end{aligned}$$

A REE of the system (2.22) and (B.3) is a stochastic process for θ_t that satisfies this system with $E_t\theta_{t+1} = \theta^e_{t+1}$. To see this possibility, note that (B.3) can be written in ARMA(1,1) form by iterating (B.3) forward by one period, and subsequently comparing this to the result one obtains by solving (B.3) for θ^e_{t+1}. This gives

$$\begin{aligned}\theta_t =& \tilde{\psi}_2^{-1}\left(\rho^{-1}\tilde{\psi}_1(1-\rho)-\tilde{\psi}_0\right)-\tilde{\psi}_1\tilde{\psi}_2^{-1}\rho^{-1}y_{t-1} \\ &+\tilde{\psi}_2^{-1}\theta_{t-1}+d_1\varepsilon_t+d_2\eta_t,\end{aligned} \tag{B.4}$$

with d_1 and d_2 being arbitrary parameters, and $\eta_t := E_t[\theta_{t+1}]-E_{t-1}[\theta_{t+1}]$ being a martingale difference sequence with $E_t[\eta_{t+1}]=0$ by the law of iterated expectations. No restrictions are imposed on d_1 or d_2, since RE formed according to (B.4) regarding θ_{t+1} are unaffected by those parameters. Therefore there is a continuum of possible solutions to (B.4). Evans and Honkapohja (1986) have shown that any finite degree ARMA solution of an equation in the form of (B.3) can at most be ARMA(1,1), and the particular form of (B.4) nests all possible ARMA solutions of finite degree. The ARMA class of solutions is stable if $|\psi_2|>1$, and is unstable for $|\psi_2|<1$.

. In this case the solution to (2.21) and (2.22) is the fundamental or minimal-state-variable (MSV) solution; it is impossible to delete any state variable from the minimum set and still obtain solutions to (2.22) and (B.3) for all permitted parameter values (McCallum, 1983). The MSV solution here is guessed to be

B.1.3 ARMA(1,1) and the MSV solution

Derivation of MSV solution: (B.4) can be re-written as

$$\theta_t = \frac{\rho\tilde{\psi}_0-\tilde{\psi}_1(1-\rho)}{\rho(1-\tilde{\psi}_2)}+\frac{\tilde{\psi}_1}{\rho(L-\tilde{\psi}_2)}y_{t-1}-\frac{d_1\tilde{\psi}_2}{(L-\tilde{\psi}_2)}\varepsilon_t-\frac{d_2}{(L-\tilde{\psi}_2)}\eta_t, \tag{B.5}$$

with L denoting the lag operator such that $Lx_t = x_{t-1}$. The parameters d_1 and d_2 can be chosen arbitrarily. In particular, to obtain the MSV solution $\theta_t = A+By_{t-1}+C\varepsilon_t$ one must first set $d_2=0$. (B.5) can be re-written as:

$$\theta_t = \frac{\rho\tilde{\psi}_0-\tilde{\psi}_1(1-\rho)}{\rho(1-\tilde{\psi}_2)}-(\rho^{-1}\tilde{\psi}_1 y_{t-1}-d_1\tilde{\psi}_2\varepsilon_t)\sum_{i=1}^{\infty}\tilde{\psi}_2^{-i}L^{i-1}.$$

$$\begin{aligned}\theta_t =& \frac{\rho\tilde{\psi}_0-\tilde{\psi}_1(1-\rho)}{\rho(1-\tilde{\psi}_2)}+\rho^{-1}\tilde{\psi}_1\tilde{\psi}_2^{-1}(1-\rho)\sum_{i=1}^{\infty}(\sum_{j=1}^{i}\rho^{-j})\tilde{\psi}_2^{-i}-\rho^{-1}\tilde{\psi}_1\tilde{\psi}_2^{-1}y_{t-1}\sum_{i=0}^{\infty}(\rho\tilde{\psi})^{-i} \\ &+\varepsilon_{t-1}(\rho^{-1}\tilde{\psi}_1\tilde{\psi}_2^{-1}\sum_{i=1}^{\infty}(\sum_{j=1}^{i}\rho^{-j}L^{i-j})\tilde{\psi}_2^{-i}+d_1\sum_{i=1}^{\infty}\tilde{\psi}_2^{-i}L^{i-1})+d_1\varepsilon_t.\end{aligned} \tag{B.6}$$

Therefore, to derive an MSV solution from a broader the class of ARMA(1,1) solutions, in which no lags of ε_t can remain, we therefore see from (B.6) that

$$d_1 = -\frac{\tilde{\psi}_1}{\rho\tilde{\psi}_2}\left(\frac{1}{\rho\tilde{\psi}_2} + (\frac{1}{\rho\tilde{\psi}_2})^2 + (\frac{1}{\rho\tilde{\psi}_2})^3 + ...\right), \tag{B.7}$$

$$= \frac{\tilde{\psi}_1}{\rho\tilde{\psi}_2(1-\rho\tilde{\psi}_2)} \qquad if \quad \tilde{\psi}_2 > \frac{1}{\rho} > 1, \tag{B.8}$$

which corresponds to the condition for stable ARMA(1,1) solutions. Otherwise, the MSV solution cannot be derived from the class of unstable ARMA(1,1) solutions, and is instead the only stable solution.

The REE values of the parameters A, B, and C are found using the method of undetermined coefficients:

$$A = \frac{\tilde{\psi}_0}{1-\tilde{\psi}_2} + \frac{\tilde{\psi}_1\tilde{\psi}_2(1-\rho)}{(1-\tilde{\psi}_2)(1-\tilde{\psi}_2\rho)},$$
$$B = \frac{\tilde{\psi}_1}{1-\tilde{\psi}_2\rho},$$
$$C = B\rho^{-1},$$

where we have assumed that $\tilde{\psi}_2 \neq 1$ and $\tilde{\psi}_2\rho \neq 1$.

B.1.4 Global convergence

Given the model discussed here has a unique equilibrium, and satisfies the assumptions of Evans and Honkapohja (1998) that guarantee global convergence, we simply apply their *Theorem 2* to the recursive learning algorithm given by (2.29) and (2.31).

For R_t, using $Ez_tz_t^{'} = M_z$, where M_z is some positive definite matrix, taking expectations we have the ODE,

$$\frac{dR}{d\tau} = M_z - R, \tag{B.9}$$

which is globally asymptotically stable and independent of x_t.

It is possible that for some t R_t may not be invertible, though this will happen only a finite number of times with probability 1. We modify the algorithm for x_t to

$$x_t = x_{t-1} + t^{-1}u(R_t)z_{t-1}\left\{z_{t-1}'\left[T\left(\hat{A}_t, \hat{B}_t\right) - x_{t-1}\right] + \eta_t\right\}, \tag{B.10}$$

where $u(R)$ is a bounded regular function from the space of $2x2$ matrices to the subspace of positive definite matrices such that $u(R) = R^{-1}$ in the neighbourhood of M_z. Then taking expectations the ODE is given by

$$\frac{dx}{d\tau} = u(R)M_z(T(\hat{A},\hat{B}) - (A,B))' \tag{B.11}$$
$$= u(R)M_z(\psi_2 - 1)((\hat{A},\hat{B}) - (A,B))'. \tag{B.12}$$

Given that the other requirements of the theorem are trivially satisfied, then it applies, and this differential equation is clearly globally asymptotically stable for $\psi_2 < 1$, and this stability is exponential; $(\hat{A},\hat{B}) \to (A,B)$ globally almost surely.

B.1.5 ALM and T-mapping ARMA solution and E-stability

$$\begin{aligned}\theta_t = &\frac{\tilde{\psi}_0 + \tilde{\psi}_2(a + b_1(1-\rho))}{1-\tilde{\psi}_2 c_1} + \frac{\tilde{\psi}_1 + \tilde{\psi}_2(b_2 + b_1\rho)}{1-\tilde{\psi}_2 c_1} y_{t-1} + \frac{\tilde{\psi}_2(b_1 + d_1) + \tilde{\psi}_1\rho^{-1}}{1-\tilde{\psi}_2 c_1}\varepsilon_t \\ &+ \frac{\tilde{\psi}_2 f_1}{1-\tilde{\psi}_2 c_1}\eta_t + \frac{\tilde{\psi}_2}{1-\tilde{\psi}_2 c_1}\sum_{j=3}^{s} b_j y_{t+1-j} + \frac{\tilde{\psi}_2}{1-\tilde{\psi}_2 c_1}\sum_{j=2}^{r} c_j \theta_{t+1-j} \\ &+ \frac{\tilde{\psi}_2}{1-\tilde{\psi}_2 c_1}\sum_{j=2}^{q} d_j \varepsilon_{t+1-j} + \frac{\tilde{\psi}_2}{1-\tilde{\psi}_2 c_1}\sum_{j=2}^{l} f_j \eta_{t+1-j}.\end{aligned} \tag{B.13}$$

This defines again a T-mapping from the PLM to the ALM with corresponding elements:

$$a = \frac{\tilde{\psi}_0 + \tilde{\psi}_2(a + b_1(1-\rho))}{1-\tilde{\psi}_2 c_1}, \tag{B.14}$$
$$b_1 = \frac{\tilde{\psi}_1 + \tilde{\psi}_2(b_1\rho + b_2)}{1-\tilde{\psi}_2 c_1}, \tag{B.15}$$
$$d_0 = \frac{\tilde{\psi}_1\rho^{-1} + \tilde{\psi}_2(b_1 + d_1)}{1-\tilde{\psi}_2 c_1}, \tag{B.16}$$
$$b_j = \frac{\tilde{\psi}_2}{1-\tilde{\psi}_2 c_1} b_{j+1}, \quad j = 2,...,s-1, \quad b_s = 0, \tag{B.17}$$
$$c_j = \frac{\tilde{\psi}_2}{1-\tilde{\psi}_2 c_1} c_{j+1}, \quad j = 1,...,r-1, \quad c_r = 0, \tag{B.18}$$
$$d_j = \frac{\tilde{\psi}_2}{1-\tilde{\psi}_2 c_1} d_{j+1}, \quad j = 1,...,q-1, \quad d_q = 0, \tag{B.19}$$
$$f_j = \frac{\tilde{\psi}_2}{1-\tilde{\psi}_2 c_1} f_{j+1}, \quad j = 0,...,l-1, \quad f_l = 0. \tag{B.20}$$

Since (B.14) - (B.20) describes a non-linear system of differential equations, we first have to linearise (B.13) to study stability properties. However, the subsystem (B.18) is independent of the other equations and can be analysed separately. The eigenvalues of the Jacobian of $T(\mathbf{c})-\mathbf{c}$ at the REE values $c_j=0$ for $j=1,...,r$ are found to be r times repeatedly equal to -1 and therefore the subsystem (B.18) will converge towards the REE values. Due to the convergence of $\mathbf{c}$ it is apparent that $\mathbf{d}$ (apart from d_0) and $\mathbf{f}$ will also converge to their REE values of vectors of zeros. Moreover, $b_j=0$ for $j=2,...,s$ is easily verified to be the values towards which the economy under learning converges. Finally, convergence of a, b_1 and d_0 are studied by analysing the Jacobian of the system (B.14)-(B.16). If this Jacobian has eigenvalues strictly less than unity, then the whole system is E-stable. It can easily be verified that the eigenvalues are ψ_2 and $\psi_2\rho$.

B.1.6 A non-steady-state PLM

The system defined as (2.38)-(2.40), linearised around steady-state values $\bar{v},\bar{n},\bar{y}=1$ has derived parameter values as follows,

$$\begin{aligned}
\kappa_0 &= (1-\kappa_2)\,\bar{\theta}-\bar{y}\kappa_1,\\
\kappa_1 &= \frac{\delta(1-\beta)(1-\bar{n})\bar{\theta}q(\bar{\theta})}{c\alpha},\\
\kappa_2 &= \delta\left[(1-\lambda)-\frac{\beta\bar{\theta}q(\bar{\theta})}{\alpha}\right] \quad (=\psi_2),\\
\kappa_3 &= \bar{\theta}\kappa_2,\\
\kappa_4 &= -\bar{\theta},\\
\phi_0 &= \alpha q(\bar{\theta})(\bar{v}+\bar{\theta}\bar{n}),\\
\phi_1 &= (1-\lambda)-\alpha\theta q(\bar{\theta}),\\
\phi_2 &= q(\bar{\theta})(1-\alpha).
\end{aligned}$$

Given the PLM (2.41), agents form expectations according to

$$v^e_{t+1}=\hat{A}_t+\hat{B}_t\left[(1-\rho)+\rho y_t\right]+\hat{C}_t n_t, \tag{B.21}$$

and the ALM is given by

$$v_t=\tilde{\kappa}_0+\tilde{\kappa}_2\left[\hat{A}_t+(1-\rho)\hat{B}_t\right]+\left[\tilde{\kappa}_1+\tilde{\kappa}_2\rho\hat{B}_t\right]y_t+\left[\tilde{\kappa}_3+\tilde{\kappa}_2\hat{C}_t\right]n_t, \tag{B.22}$$

where

$$\tilde{\kappa}_0 = \frac{\kappa_0 + \kappa_1(1-\rho) + \kappa_3\phi_0}{1-\kappa_3\phi_2},$$
$$\tilde{\kappa}_1 = \frac{\kappa_1\rho}{1-\kappa_3\phi_2},$$
$$\tilde{\kappa}_2 = \frac{\kappa_2}{1-\kappa_3\phi_2},$$
$$\tilde{\kappa}_3 = \frac{\kappa_4 + \kappa_3\phi_1}{1-\kappa_3\phi_2}.$$

Given the mapping $\tilde{T}$ defined in the main text, the REE is E-stable if all the eigenvalues of the Jacobian of $\tilde{T}(\hat{A},\hat{B},\hat{C}) - (\hat{A},\hat{B},\hat{C})$ have negative real parts. Thus, we must have

$$\tilde{\kappa}_2\rho - 1 < 0$$

and

$$\tilde{\kappa}_2 - 1 < 0,$$

whereby the second condition implies the validity of the first. Therefore, we need to check for what range of parameter values of the model the second condition is true. Writing out the term $\tilde{\kappa}_2$ and rearranging, we see that the required condition is

$$\delta\left[(1-\lambda) - \frac{\beta\bar{\theta}q(\bar{\theta})}{\alpha}\right]\left[1+\bar{\theta}q(\bar{\theta})(1-\alpha)\right] < 1, \tag{B.23}$$

or

$$\psi_2\left[1+\bar{\theta}q(\bar{\theta})(1-\alpha)\right] < 1. \tag{B.24}$$

Given that $\tilde{\kappa}_2 \geq \psi_2$, if the E-stability condition holds with this alternative PLM, then convergence will be slower. For the complete range of possible model parameters, this condition does not hold. As realistic levels of λ are small, the condition would be sensitive to assumed parameter values of β and α. For example, given $\alpha > \beta > 0$, which is the case of low worker bargaining power, whereby wages are reduced towards the value of the outside option, and there is excessive firm entry, or inefficiently high according to the Hosios (1990) condition, it is more likely E-stability will not hold.

Appendix B.2 Parametrisation of the model

We normalise average productivity to be one. For the productivity process we estimate an AR(1) in log deviations from trend output per worker, dynamically de-trended using the HP filter with standard quarterly smoothing parameter, and find an auto-regressive parameter ρ for the period of 0.84, and a standard deviation for the shocks σ_ε of 0.0063 (assuming them to be normally distributed). For the labour market, we parametrise the model to the unemployment rate, measured as the fraction of the economically active population aged 16 and over who are ILO unemployed. We use official quarterly time series from Office for National Statistics (ONS) Labour Market Statistics. For transition rates between labour market states we use the flows time series similarly published by ONS, which are derived from the Two-quarter Longitudinal Labour Force Survey and are consistent with all stocks series. The economy we describe has two states. In reality there is a third: economic inactivity. To adhere to our interpretation of u_t as the unemployment rate, abstracting from the relative size of the inactive population over the business cycle, as is common in the literature (Shimer, 2005; Hagedorn and Manovskii, 2008), we must carefully construct from the raw data measures of job finding and separation rates. In the notation of the model, the steady-state unemployment rate is given by

$$u_t^* = \frac{\lambda}{\lambda + \theta_t q(\theta_t)}. \tag{B.25}$$

As per Smith (2011), using three-state flows data between the stocks in employment, unemployment and inactivity, denoted by $\{E, U, I\}$, with transition rates, for example between inactivity and unemployment, denoted by p_{IU_t}, we can re-write (B.25) as

$$u_t^* = \frac{p_{EU_t} + \frac{p_{EI_t} p_{IU_t}}{p_{IU_t} + p_{IE_t}}}{\underbrace{p_{EU_t} + \frac{p_{EI_t} p_{IU_t}}{p_{IU_t} + p_{IE_t}}}_{\lambda_t} + \underbrace{p_{UE_t} + \frac{p_{UI_t} p_{IE_t}}{p_{IU_t} + p_{IE_t}}}_{\theta_t q(\theta_t)}}. \tag{B.26}$$

As such, the separation rate from real data which is consistent with the model described here is the sum of the direct transition rate from employment to unemployment and a term which captures the indirect role of transitions to unemployment via inactivity - with a similar interpretation for the job finding rate.

Using this measure of the hiring rate from the transition rates data, we estimate the parameters of the aggregate matching function using least squares as follows for

2002q1-13q2:

$$\log\left(p_{UE_t} + \frac{p_{UI_t} p_{IE_t}}{p_{IU_t} + p_{IE_t}}\right) = \log(\mu) + (1-\alpha)\log\left(\frac{v_t}{u_t}\right) + \zeta_t, \tag{B.27}$$

where data for v_t come from the quarterly ONS aggregate vacancies series, and u_t is the UK national unemployment rate. Following Borowczyk-Martins et al. (2013), we consider time trends in the estimation to account for the endogeneity of unobserved shifts in the matching efficiency with the number of vacancies that firms open, but these all drop out. We also carry out tests that the matching function is Cobb-Douglas, and reject the alternative. In line with the existing literature, we find that the data suggests the matching function has decreasing returns to scale, although we proceed as though it is constant (see Pissarides and Petrongolo (2001) for a thorough review of estimates of the aggregate matching function). We find estimates of $\alpha = 0.67$ and $\mu = 0.56$. For the constant separation rate parameter in the model, over the same period we choose an average value of the two-quarter composite hazard rate: $p_{EU} + \frac{p_{EI} p_{IU}}{p_{IU} + p_{IE}} = \lambda = 0.023$. (In practice we regress the data on a constant and cubic trend to account for low frequency shifts for the short period in question, then selecting the estimated constant as the parameter value - we similarly do this when estimating moments of the labour market variables presented in Table 2.2). The discount factor is set as $\delta = 0.99$, and to restrict the number of free parameters we let the bargaining power adhere to the Hosios (1990) condition, $\beta = \alpha = 0.67$. We set the flow value of unemployment to 0.8. How to select or estimate appropriate values of both the bargaining power and the flow value of unemployment are open to debate. Shimer (2005) and subsequently Hagedorn and Manovskii (2008) are often considered in the literature as more extreme examples for parametrisations, and highlight how this affects the ability of the model to match the observed volatility of unemployment and vacancy creation. With the relatively arbitrary parametrisation applied here, we are somewhere in between these two examples. The remaining parameter, the flow vacancy cost c, is chosen to match the observed level of average labour market tightness over the period, as displayed in Table 2.2 and as used to estimate the parameters of the matching function.

Appendix B.3 Additional figures

FIGURE B1: Simulation of the equilibrium of the stochastic model with the assumption of rational expectations: the REE

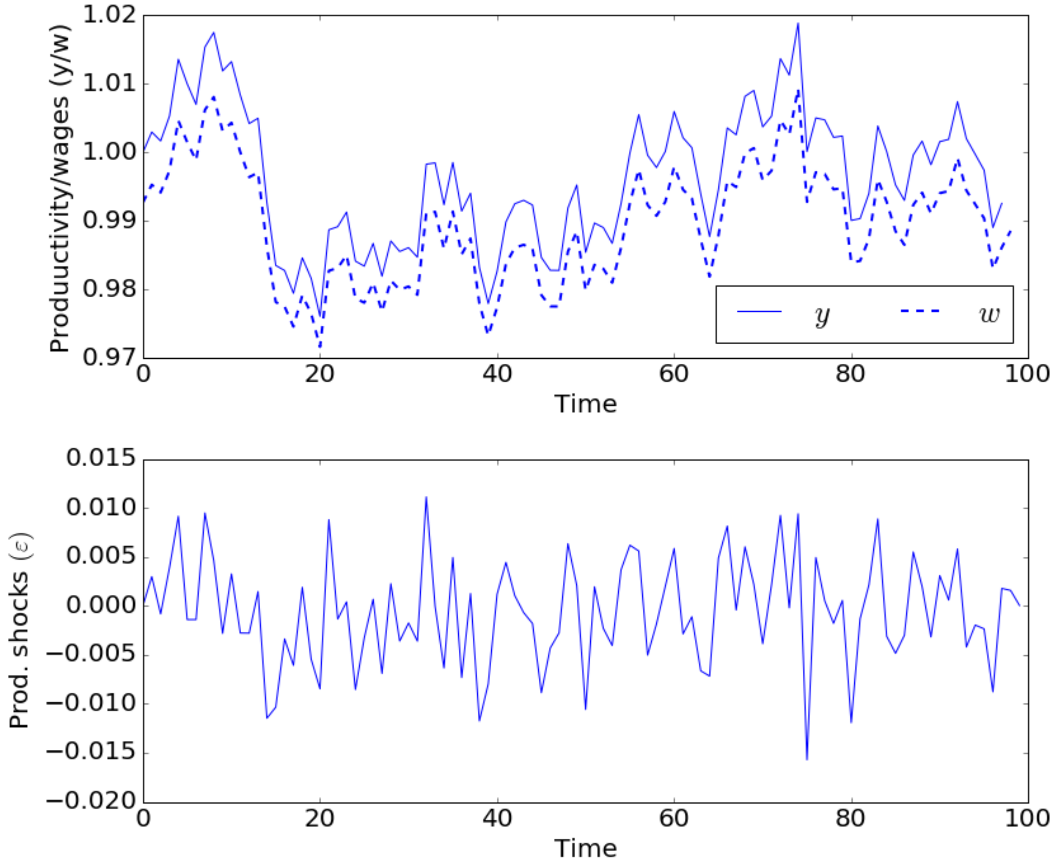

Note.- The simulation here is identical to that described under decreasing gain learning for Figure 2.2.

FIGURE B2: Simulation of the equilibrium of the stochastic model with the assumption of rational expectations: the REE

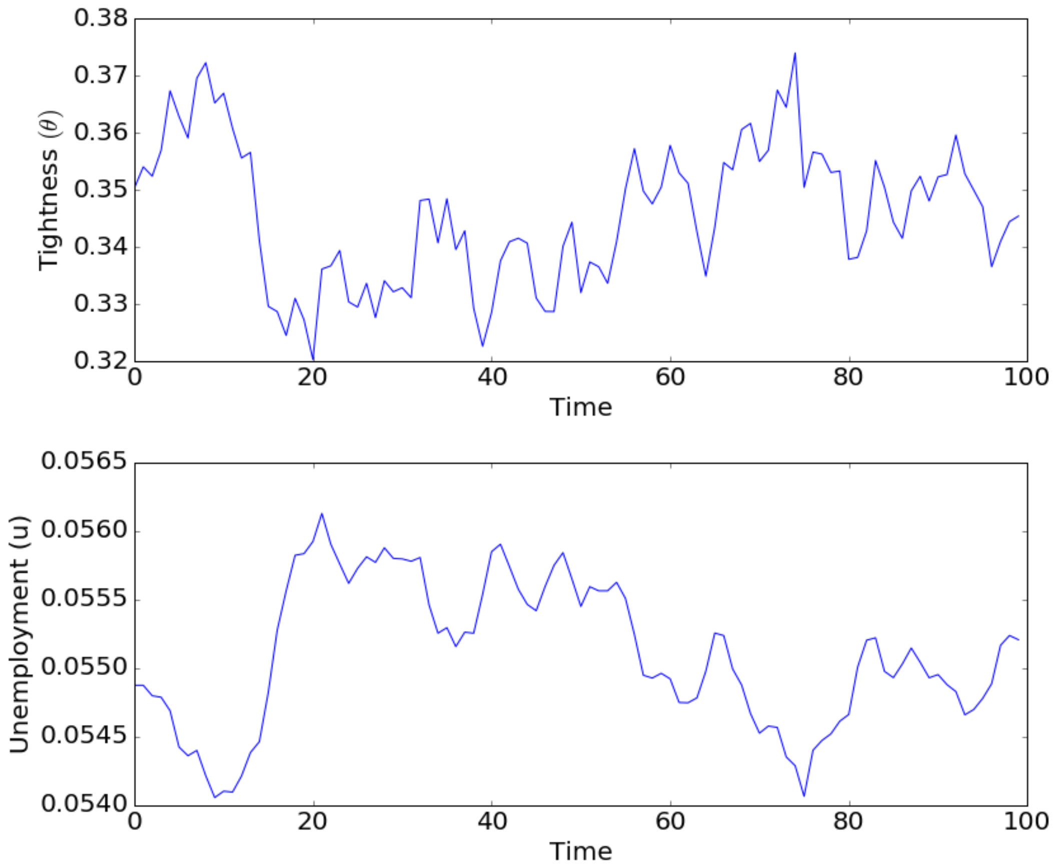

Note.- The simulation here is identical to that described under decreasing gain learning for Figure 2.2.

FIGURE B3: Convergence of agents' parameter estimates under decreasing gain learning to the REE values

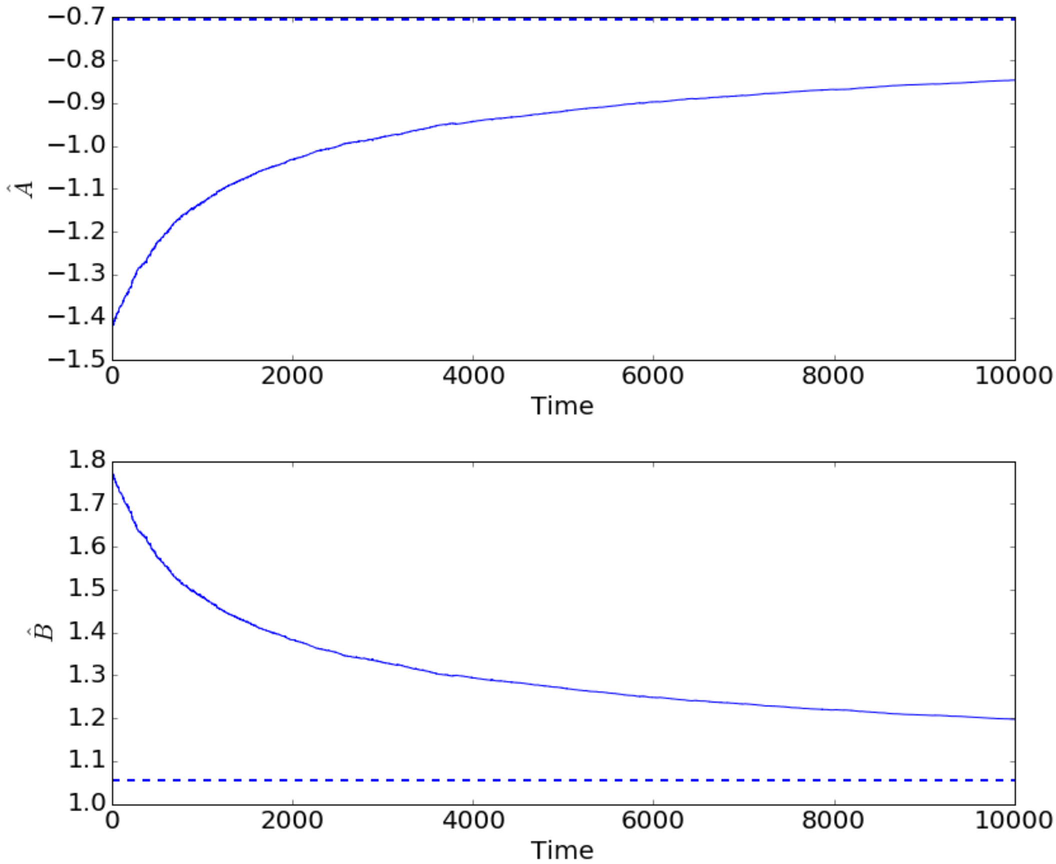

Note.- The simulation here is identical to that described under decreasing gain learning for Figure 2.2. Dashed lines give the true REE parameter values.

FIGURE B4: Comparison of sample paths for endogenous variables under RE and decreasing gain learning, and agents' parameter estimates, following a structural shock

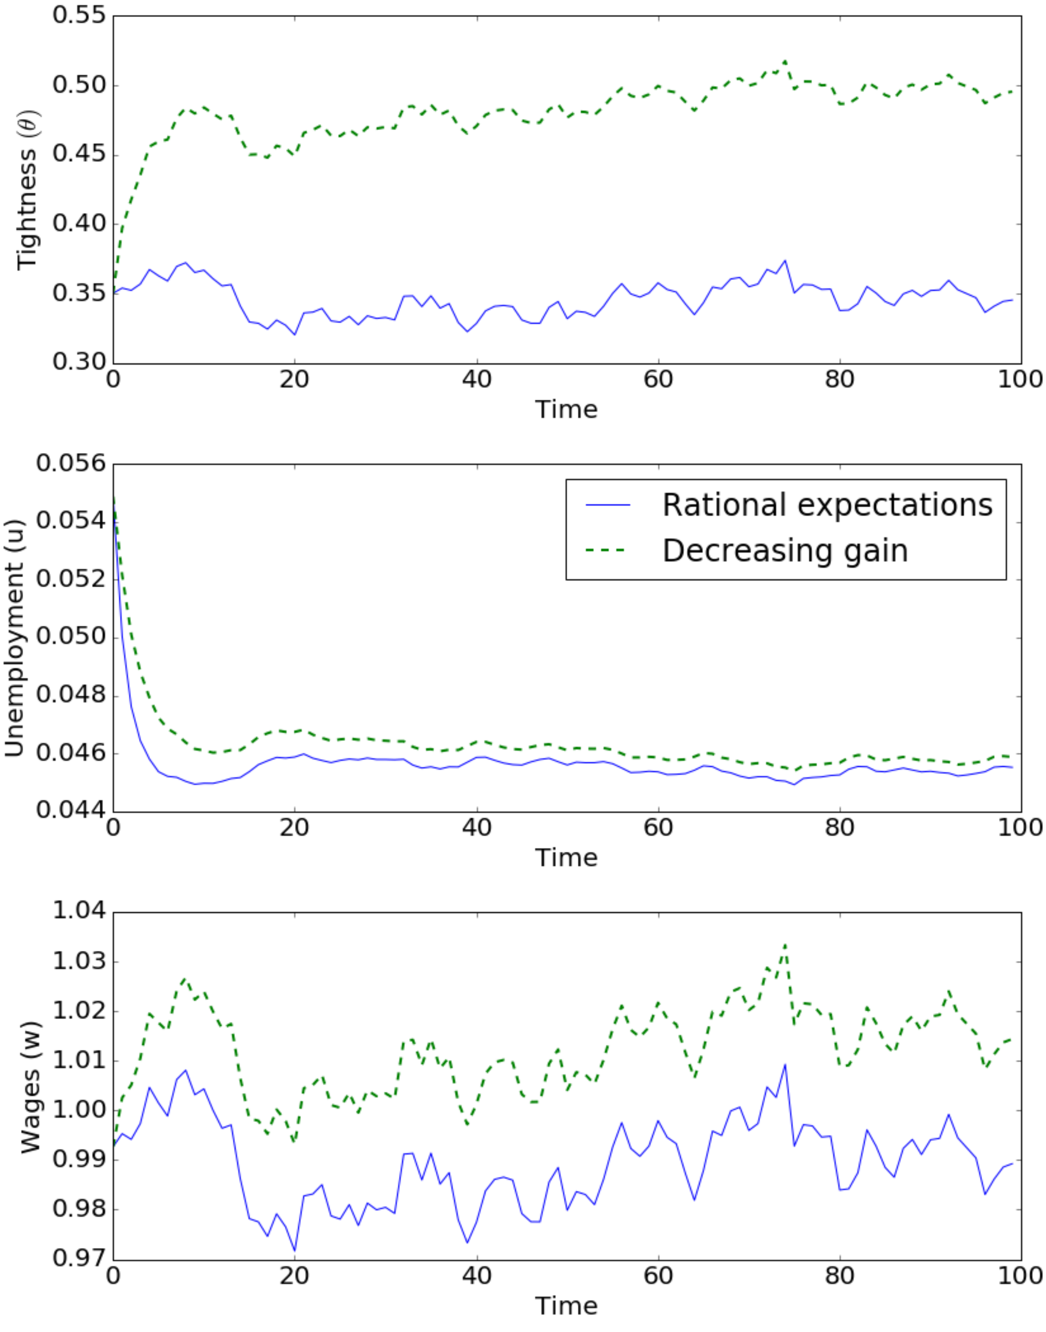

Note.- these simulation paths are the results of a negative 20% shock to the flow value of unemployment b, with initial parameter estimates assumed to be at the true pre-shock REE values.

FIGURE B5: Comparison of sample paths for endogenous variables under RE and decreasing gain learning, and agents' parameter estimates, following a structural shock

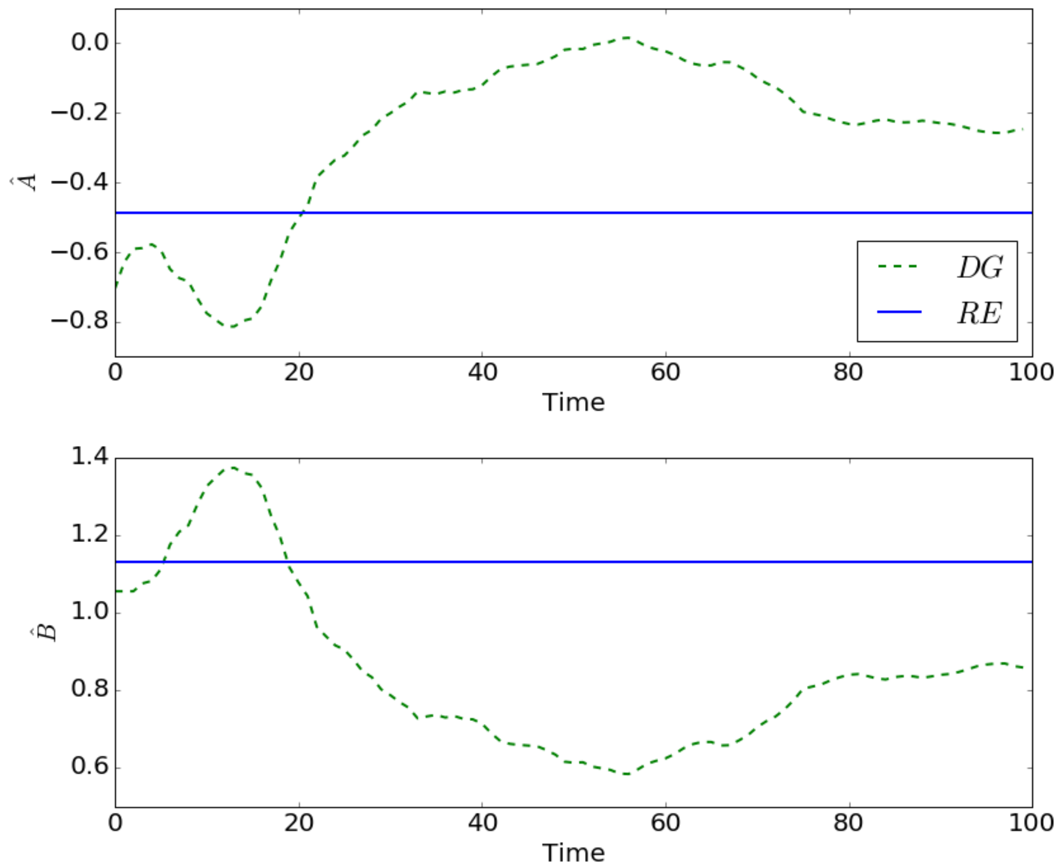

Note.- these simulation paths are the results of a negative 20% shock to the flow value of unemployment b, with initial parameter estimates assumed to be at the true pre-shock REE values.

FIGURE B6: Comparing the speed of convergence to the REE under decreasing gain learning: changing worker bargaining power β

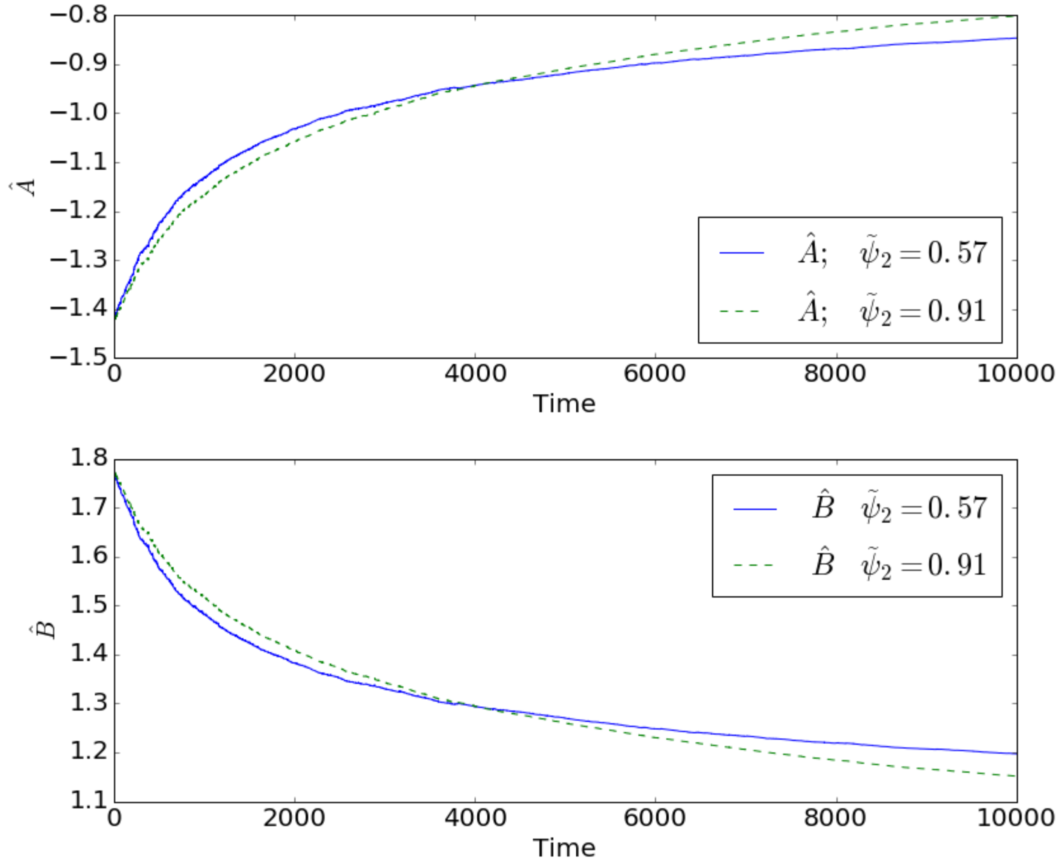

Note.- given $\beta = 0.1$, then $\psi_2 = 0.91$. For $\beta = \alpha$, $\psi_2 = 0.57$, as in Figure 2.2. The crossing of the time paths indicates a decreased speed of convergence since, for example, the REE parameter value with $\beta = 0.1$ of $B = 0.56$ is substantially lower than value with $\beta = \alpha$.

Appendix C

Measuring sectoral income shares: Accounting for input-output structures across countries

Appendix C.1 Data sources for Net Operating Surplus

This section provides information on the sources of the Net Operating Surplus (NOPS) data at the industry level. Further, the assumptions made in cases of missing data for individual industries or years on a country basis are stated. The data is obtained from two main sources: the OECD's Database for Structural Analysis (STAN) and National Accounts Aggregates by Industry from Eurostat.[1] For some countries for which no data on NOPS is available in these two databases, data from national statistics agencies have been collected.

STAN is based on OECD countries' Annual National Accounts By Activity tables and uses national industrial surveys and census data to estimate any missing detail. The latest version of STAN is based on the International Standard Industrial Classification of all economic activities, revision 4 (ISIC Rev. 4). However, due to the larger coverage of countries in the previous version, STAN based on ISIC Rev. 3 will be used, which was last updated in May 2011. Since in the WIOD SEA industries are classified according to the Statistical Classification of Economic Activities in the European Community (NACE) revision 1, NOPS industry data is linked between STAN and the WIOD using correspondence tables from the United Nations Statistics Division. All data in national currencies has been converted into U.S. Dollar using the exchange rates

[1] Source OECD's STAN: https://stats.oecd.org/...STAN08BIS; accessed 04/12/14. Source Eurostat: http://appsso.eurostat.ec.europa.eu/...lang=en; accessed 04/12/14.

provided by the WIOD. Data on NOPS for the following countries is available from OECD STAN: Austria, Czech Republic, Estonia, Finland, Hungary, Luxembourg, Netherlands, and Slovenia. STAN data, combined with imputations which are listed below, is also used for Belgium, Germany, Denmark, Korea, and Sweden.

Eurostat collects data from EU member states and also compiles data at the industry level. The major data source for this are enterprise surveys, production surveys and annual reports or business accounts from major companies. Eurostat's National Accounts Aggregates by Industry Database uses a different classification system than the WIOD, namely NACE Rev. 2, and thus correspondence tables to NACE Rev. 1, available from Eurostat, are used to allocate NOPS to industries in the WIOD. The aggregate "Manufacture of textiles, wearing apparel, leather and related products" needs to be split between NACE Rev. 1 codes 17t18 and 19. For this I use value added shares from the WIOD. Complete data is available for Bulgaria, Cyprus, France, Greece, Italy, Ireland, Poland, Portugal, Romania, and the United Kingdom. Imputations, explained below, were made for Latvia, Malta, and Spain.

Belgium, Denmark, Germany, and Sweden: NOPS for 2009 is imputed using least squares projections on value added from simple regressions of NOPS on value added for each individual industry between 1995-2008. Additionally for Sweden, NOPS data on industries 60-63 is only available as aggregate prior to 2004. I allocate this sum to industries by applying the 2004 shares in the following years. **Korea**: No disaggregated data is available on sub-industries in trade (NACE 50-52) and transport (NACE 60-63). I allocate these aggregates according to value added shares from the WIOD. **Latvia, and Spain**: NOPS at the industry level for the period 1995-1999 is imputed by applying the shares in total NOPS from 2000 in the previous years. **Malta**: Only the aggregate value of NOPS for industries 62 (Air transport) and 64 (Postal services and telecommunication) is available. This is allocated according to value added shares from the WIOD. No data is available for category 23 (Coke and refined petroleum products), so I set NOPS equal to value added in the WIOD.

Australia: NOPS data is obtained from the Australian System of National Accounts - 5204.0, Tables 46 & 47.[2] Industries between ANZSIC and the WIOD are mapped according to the EU KLEMS correspondence table. Aggregate value of "transport" is allocated to NACE 60-63 according to value added shares from the WIOD. **Russia**: Agricultural labour income is based on Mincer-type regressions, and for other industries, wages of employees are imputed for the self-employed (Voskoboynikov, 2012). Data on NOPS is generally not available, thus Gross

[2] Available at http://www.abs.gov.au/...Document; accessed 01/06/15.

Operating Surplus (GOPS), obtained from Rosstat, (only post 2001) for main sectors is used instead. I allocate GOPS according to value added shares from the WIOD. Prior to 2001 I use least squares projections on value added from the WIOD. **USA**: The GDP-by-Industry accounts of the Bureau of Economic Analysis (BEA) contain information on mixed income at the industry-level for 1998-2009. Data for the years 1995-1997 is projected on value added from the WIOD by regressing mixed income on industry level value added for 1998-2009. Further, only the aggregate of NACE 50 and 51 is available. This value is allocated to its sub-categories using value added shares from the WIOD.

For certain countries, no adjustments were made, because the data on labour compensation from the WIOD is based on micro-level regressions which take into account observable characteristics of the self-employed, obtained from additional non-public data sources. These countries are **Brazil, Canada, China, Indonesia, India, Japan, and Taiwan**.

Appendix C.2 Industry and commodity classifications

Expenditure on products (goods as well as services) in the WIOD is classified according to the 2008 version of the *statistical classification of products by activity* (CPA) in the European Economic Community.[3] The sub-categories $j \in G$ which belong to the **goods sector** are: Products of agriculture, hunting and related services (1), Products of forestry, logging and related services (2), Fish and other fishing products; services incidental of fishing (5), Coal and lignite; peat (10), Crude petroleum and natural gas; services incidental to oil and gas extraction excluding surveying (11), Uranium and thorium ores (12), Metal ores (13), Other mining and quarrying products (14), Food products and beverages (15), Tobacco products (16), Textiles (17), Wearing apparel; furs (18), Leather and leather products (19), Wood and products of wood and cork (except furniture); articles of straw and plaiting materials (20), Pulp, paper and paper products (21), Printed matter and recorded media (22), Coke, refined petroleum products and nuclear fuels (23), Chemicals, chemical products and man-made fibres (24), Rubber and plastic products (25), Other non-metallic mineral products (26), Basic metals (27), Fabricated metal products, except machinery and equipment (28), Machinery and equipment n.e.c. (29), Office machinery and computers (30), Electrical machinery and apparatus n.e.c. (31), Radio, television and communication equipment and apparatus (32), Medical, precision and optical instruments, watches and clocks (33), Motor vehicles, trailers and semi-trailers (34), Other transport equipment (35), Furniture; other manufactured goods n.e.c.

[3] Source: http://ec.europa.eu/.../cpa-2008; accessed 13/04/16.

(36), Secondary raw materials (37), Electrical energy, gas, steam and hot water (40), Collected and purified water, distribution services of water (41), and Construction work (45).

The sub-categories $j \in S$ which belong to the **services sector** are: Trade, maintenance and repair services of motor vehicles and motorcycles; retail sale of automotive fuel (50), Wholesale trade and commission trade services, except of motor vehicles and motorcycles (51), Retail trade services, except of motor vehicles and motorcycles; repair services of personal and household goods (52), Hotel and restaurant services (55), Land transport; transport via pipeline services (60), Water transport services (61), Air transport services (62), Supporting and auxiliary transport services; travel agency services (63), Post and telecommunication services (64), Financial intermediation services, except insurance and pension funding services (65), Insurance and pension funding services, except compulsory social security services (66), Services auxiliary to financial intermediation (67), Real estate services (70), Renting services of machinery and equipment without operator and of personal and household goods (71), Computer and related services (72), Research and development services (73), Other business services (74), Public administration and defence services; compulsory social security services (75), Education services (80), Health and social work services (85), Sewage and refuse disposal services, sanitation and similar services (90), Membership organisation services n.e.c. (91), Recreational, cultural and sporting services (92), Other services (93), Private households with employed persons (95).

TABLE C.1: Classification of industries in the WIOD into the goods and services sector

Goods		Services	
Industry	Code	Industry	Code
Agr., forestry & fishing	AtB	Motor veh. & fuel trade	50
Mining & quarrying	C	Wholesale trade	51
Food, bever. & tobacco	15t16	Retail trade	52
Textiles	17t18	Hotels & restaurants	H
Leather & footwear	19	Land transport	60
Wood products	20	Water transport	61
Paper, printing & publ.	21t22	Air transport	62
Coke & refined petrol.	23	Transport services	63
Chemical products	24	Post & telecomm.	64
Rubber & plastics	25	Financial services	J
Non-metal. mineral prod.	26	Real estate	70
Basic & fabric. metal	27t28	Business services	71t74
Other machinery	29	Government	L
Electr. & optical equip.	30t33	Education	M
Transport equip.	34t35	Health	N
Other manufacturing	36t37	Other services	O
Utilities	E	Households w/ empl. pers.	P
Construction	F		

Notes.- industries included in the goods and services sector. The codes in the WIOD follow the International Standard Industrial Classification of All Economic Activities (ISIC) Rev. 3 of the United Nations. Source: http://unstats.un.org/.../regcst.asp?Cl=2; accessed 13/04/16.

TABLE C.2: Allocation to the goods and services categories of the World Bank's ICP 2005

Sector	Description	Code
Goods sector	Food and non-alcoholic beverages	1101
	Alcoholic beverages and tobacco	1102
	Clothing and footwear	1103
	Housing, water, electricity,	1104
	gas, and other fuels	1104
	Furnishings, household equipment, and household maintenance	1105
	Miscellaneous goods and services	1112
	Machinery and equipment	1501
	Construction	1502
	Other products	1503
Services sector	Health	1106
	Transport	1107
	Communication	1108
	Recreation and culture	1109
	Education	1110
	Restaurant and hotels	1111

Notes.- descriptions and codes of the World Bank's ICP 2005. Categories and their associated expenditure sector in the WIOD.

Appendix C.3 Decomposition: derivations

The difference between the labour income shares of country i and the US in sector $z \in \{G, S\}$ for a given year is

$$\alpha_{iz} - \alpha_{USz} = \mathbf{r}_i' \mathbf{L}_i \mathbf{e}_{iz} \left[\mathbf{V}_i' \mathbf{L}_i \mathbf{e}_{iz}\right]^{-1} - \mathbf{r}_{US}' \mathbf{L}_{US} \mathbf{e}_{USz} \left[\mathbf{V}_{US}' \mathbf{L}_{US} \mathbf{e}_{USz}\right]^{-1}. \tag{C.1}$$

The difference operator Δ is defined as $\Delta \mathbf{a} = \mathbf{a}_i - \mathbf{a}_{US}$ for any vector or matrix $\mathbf{a}$. Adding and subtracting terms gives

$$\begin{aligned}
(\alpha_{iz} - \alpha_{USz}) \left[\mathbf{V}_i' \mathbf{L}_i \mathbf{e}_{iz}\right] \left[\mathbf{V}_{US}' \mathbf{L}_{US} \mathbf{e}_{USz}\right] = \\
&\mathbf{r}_i' \mathbf{L}_i \mathbf{e}_{iz} \left[\mathbf{V}_{US}' \mathbf{L}_{US} \mathbf{e}_{USz}\right] - \mathbf{r}_{US}' \mathbf{L}_{US} \mathbf{e}_{USz} \left[\mathbf{V}_i' \mathbf{L}_i \mathbf{e}_{iz}\right] \\
&= (\Delta \mathbf{r}_i)' \mathbf{L}_i \mathbf{e}_{iz} \left[\mathbf{V}_{US}' \mathbf{L}_{US} \mathbf{e}_{USz}\right] + \mathbf{r}_{US}' (\Delta \mathbf{L}_i) \mathbf{e}_{iz} \left[\mathbf{V}_{US}' \mathbf{L}_{US} \mathbf{e}_{USz}\right] \\
&+ \mathbf{r}_{US}' \mathbf{L}_{US} (\Delta \mathbf{e}_{ij}) \left[\mathbf{V}_{US}' \mathbf{L}_{US} \mathbf{e}_{USz}\right] - (\Delta \mathbf{V}_i)' \mathbf{L}_i \mathbf{e}_{iz} \left[\mathbf{r}_{US}' \mathbf{L}_{US} \mathbf{e}_{USz}\right] \\
&- \mathbf{V}_{US}' (\Delta \mathbf{L}_i) \mathbf{e}_{iz} \left[\mathbf{r}_{US}' \mathbf{L}_{US} \mathbf{e}_{USz}\right] - \mathbf{V}_{US}' \mathbf{L}_{US} (\Delta \mathbf{e}_{iz}) \left[\mathbf{r}_{US}' \mathbf{L}_{US} \mathbf{e}_{USz}\right].
\end{aligned}$$

Gathering terms gives

$$\alpha_{iz} - \alpha_{USz} = \underbrace{(\Delta \mathbf{S}_{iz})' \mathbf{L}_i \mathbf{e}_{iz}}_{\text{VA labour share}} + \underbrace{\mathbf{S}'_{iz}(\Delta \mathbf{L}_i)\mathbf{e}_{iz}}_{\text{Supply chain}} + \underbrace{\mathbf{S}'_{iz}\mathbf{L}_{US}(\Delta \mathbf{e}_{iz})}_{\text{Expenditure weights}} \quad . \tag{C.2}$$

The vector $(\Delta \mathbf{S}_{iz})'$ consists of the differences between industry level labour income and value added in country i and the US, and weighting terms (see below). By post multiplication of this vector with country i's supply chain network and expenditure weights, those differences are translated into changes of sectoral labour income shares at the expenditure side and hence the effect of varying industry labour income shares, while keeping all other things equal, can be measured in this way.

Similarly, $(\Delta \mathbf{L}_i)$ is a matrix which captures the deviations in the supply chain networks of country i and the US. Entries of this matrix equal zero whenever country i and the US exhibit the same input-output linkages. For example, if a larger fraction of the total supply of a certain commodity is produced by a particular industry in country i than in the US, the entry of $(\Delta \mathbf{L}_i)$ representing this industry-commodity combination will be positive and equal to the difference in supply shares.

The last term refers to deviations of the expenditure vectors between country i and the US, which means that it measures the effect that aggregation has on the sectoral labour income shares. Each of the components in the above equation gives the observed difference in labour income shares, if the other two components were identical to the US values.

The vector that captures the share of labour income in value added is

$$\mathbf{S}'_{iz} = \mathbf{r}'_{US} X_{iz} - \mathbf{V}'_{US} Y_{iz} \tag{C.3}$$

and, with a slight abuse of notation, the disparity between country i and the corresponding US values is

$$(\Delta \mathbf{S}_{iz})' = (\Delta \mathbf{r}_i)' X_{iz} - (\Delta \mathbf{V}_i)' Y_{iz}. \tag{C.4}$$

The country and sector specific weights, X_{iz} and Y_{iz}, are

$$X_{iz} = \mathbf{V}'_{US}\mathbf{L}_{US}\mathbf{e}_{USz}\left[\mathbf{V}'_i\mathbf{L}_i\mathbf{e}_{iz}\mathbf{V}'_{US}\mathbf{L}_{US}\mathbf{e}_{USz}\right]^{-1} \tag{C.5}$$

$$Y_{iz} = \mathbf{r}'_{US}\mathbf{L}_{US}\mathbf{e}_{USz}\left[\mathbf{V}'_i\mathbf{L}_i\mathbf{e}_{iz}\mathbf{V}'_{US}\mathbf{L}_{US}\mathbf{e}_{USz}\right]^{-1}. \tag{C.6}$$

The weights X_{iz} and Y_{iz} normalise the quantities of labour and total compensation by the relative size of the compared economies, which is necessary since these values are not ratios. Note that varying labour income while holding capital income constant, thus allowing value added to change, is an alternative of allocating the difference in labour income across countries to the terms that form value added. Alternatively, one could treat value added as fixed and vary only labour income while adjusting capital income in an offsetting way to guarantee value added equals labour income plus capital income. However, the latter method has an ambiguous interpretation as it only captures differences in labour compensation and not in the labour income *share*. Once the above terms are derived for each year during 1995-2009, the arithmetic average is computed to arrive at the result of the decomposition shown in the main text.

Appendix C.4 Additional tables and figures

TABLE C.3: Labour income shares, averages over 1995-2009

	Aggregate economy	Goods sector	Services sector
AUS	0.547	0.488	0.584
AUT	0.586	0.574	0.594
BEL	0.598	0.592	0.602
BGR	0.460	0.451	0.468
BRA	0.520	0.480	0.545
CAN	0.579	0.547	0.602
CHN	0.484	0.491	0.471
CYP	0.541	0.512	0.551
CZE	0.509	0.504	0.515
DEU	0.596	0.638	0.566
DNK	0.637	0.609	0.652
ESP	0.578	0.574	0.581
EST	0.552	0.566	0.542
FIN	0.585	0.572	0.594
FRA	0.591	0.611	0.581
GBR	0.616	0.585	0.630
GRC	0.490	0.508	0.485
HUN	0.548	0.535	0.558
IDN	0.486	0.420	0.606
IND	0.527	0.515	0.543
IRL	0.527	0.461	0.590
ITA	0.522	0.528	0.519
JPN	0.507	0.532	0.492
KOR	0.557	0.543	0.570
LTU	0.478	0.466	0.486
LUX	0.534	0.600	0.516
LVA	0.497	0.496	0.497
MLT	0.527	0.479	0.548
NLD	0.589	0.549	0.611
POL	0.484	0.466	0.500
PRT	0.599	0.597	0.600
ROU	0.497	0.533	0.458
RUS	0.504	0.493	0.515
SVN	0.618	0.625	0.611
SWE	0.627	0.609	0.637
TWN	0.599	0.567	0.625
USA	0.566	0.574	0.563

Notes.- labour income shares are the arithmetic average of the annual shares over 1995-2009. Annual shares are computed using the method described in section 3.2. Labour income of the self-employed is imputed assuming equal hourly wage rates between employees and the self-employed, constrained by imputed OSPUE.

Bibliography

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

Acemoglu, D., and V. Guerrieri. 2008. "Capital Deepening and Nonbalanced Economic Growth." *Journal of Political Economy*, 116(3): 467–498.

Andolfatto, D. 1996. "Business Cycles and Labor-Market Search." *American Economic Review*, 86(1): 112–132.

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

Atkinson, A. B. 2009. "Factor shares: the principal problem of political economy?" *Oxford Review of Economic Policy*, 25(1): 3–16.

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

Azmat, G., A. Manning, and J. Van Reenen. 2012. "Privatization and the Decline of Labour's Share: International Evidence from Network Industries." *Economica*, 79 470–492.

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

Barro, R. J., and J.-W. Lee. 2013. "A new data set of educational attainment in the world, 1950-2010." *Journal of Development Economics*, 104(Supplement C): 184–198.

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

Benhabib, J., and R. E. Farmer. 1999. "Indeterminacy and Sunspots in Macroeconomics." *Handbook of Economic Growth, Volume 2:* 387–448.

Benveniste, A., M. Métivier, and P. Priouret. 2012. *Adaptive algorithms and stochastic approximations*. 22 Springer Science & Business Media.

Berardi, M., and J. K. Galimberti. 2012. "On the plausibility of adaptive learning in macroeconomics: A puzzling conflict in the choice of the representative algorithm." Centre for Growth and Business Cycle Research Discussion Paper Series 177, The Univeristy of Manchester.

Berardi, M., and J. K. Galimberti. 2013. "A note on exact correspondences between adaptive learning algorithms and the Kalman filter." *Economics Letters*, 118(1): 139–142.

Bernanke, B. S., and R. S. Gürkaynak. 2001. "Is Growth Exogenous? Taking Mankiw, Romer, and Weil Seriously." *NBER Macroeconomics Annual*, 16: 11–57.

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

Blanchard, O., and M. W. Watson. 1983. "Bubbles, Rational Expectations and Financial Markets." NBER Working Papers 0945, National Bureau of Economic Research.

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

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

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

Borowczyk-Martins, D., G. Jolivet, and F. Postel-Vinay. 2013. "Accounting For Endogeneity in Matching Function Estimation." *Review of Economic Dynamics*, 16(3): 440–451.

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

Branch, W., and B. McGough. 2016. "Heterogeneous beliefs and trading inefficiencies." *Journal of Economic Theory*, 163(C): 786–818.

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

Bullard, J., and K. Mitra. 2002. "Learning about monetary policy rules." *Journal of Monetary Economics*, 49(6): 1105–1129.

Burgess, S., and H. Turon. 2010. "Worker flows, job flows and unemployment in a matching model." *European Economic Review*, 54(3): 393–408.

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

Carceles-Poveda, E., and C. Giannitsarou. 2007. "Adaptive learning in practice." *Journal of Economic Dynamics and Control*, 31(8): 2659–2697.

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

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

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

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

Di Pace, F., K. Mitra, and S. Zhang. 2016. "Adaptive learning and labour market dynamics." Bank of England working papers 633, Bank of England.

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

DiNardo, J., N. M. Fortin, and T. Lemieux. 1996. "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach." *Econometrica*, 64(5): 1001–1044.

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

Elsby, M. W. L., B. Hobijn, and A. Sahin. 2013. "The decline of the U.S. labor share." *Brookings Papers on Economic Activity*, Fall(2): 1–52.

Elsby, M. W. L., D. Shin, and G. Solon. 2016. "Wage Adjustment in the Great Recession and Other Downturns: Evidence from the United States and Great Britain." *Journal of Labor Economics*, 34(S1): S249 – S291.

Elsby, M. W. L., and J. C. Smith. 2010. "The Great Recession in the UK labour market: A transatlantic perspective." *National Institute Economic Review*, 214 R1–R12.

Eusepi, S., and B. Preston. 2011. "Expectations, Learning, and Business Cycle Fluctuations." *American Economic Review*, 101(6): 2844–2872.

Evans, G. W., and S. Honkapohja. 1986. "A Complete Characterization of ARMA Solutions to Linear Rational Expectations Models." *Review of Economic Studies*, 53(2): 227–239.

Evans, G. W., and S. Honkapohja. 1998. "Economic Dynamics with Learning: New Stability Results." *Review of Economic Studies*, 65(1): 23–44.

Evans, G. W., and S. Honkapohja. 2001. *Learning and expectations in macroeconomics*. Frontiers of economic research, Princeton University Press, Princeton, NJ.

Feenstra, R. C., R. Inklaar, and M. P. Timmer. 2015. "The Next Generation of the Penn World Table." *American Economic Review*, 105(10): 3150–3182.

Gertler, M., C. Huckfeldt, and A. Trigari. 2016. "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires." CEPR Discussion Papers 11318, C.E.P.R. Discussion Papers.

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

Gollin, D. 2002. "Getting Income Shares Right." *Journal of Political Economy*, 110(2): 458–474.

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

Grobovšek, J. 2017. "Development accounting with intermediate goods." *B.E. Journal of Macroeconomics*, forthcoming.

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

Hagedorn, M., and I. Manovskii. 2008. "The cyclical behavior of equilibrium unemployment and vacancies revisited." *American Economic Review*, 98(4): 1692–1706.

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

Herrendorf, B., and A. Valentinyi. 2012. "Which Sectors Make Poor Countries So Unproductive?" *Journal of the European Economic Association*, 10(2): 323–341.

Herrendorf, B., R. Rogerson, and A. Valentinyi. 2014. "Growth and structural transformation." *Handbook of Economic Growth, Volume 2:* 855–941.

Horowitz, K. J., and M. A. Planting. 2006. "Concepts and Methods of the U.S. Input-Output Accounts." BEA Papers 0066, Bureau of Economic Analysis.

Hosios, A. 1990. "On The Efficiency of Matching and Related Models of Search and Unemployment." *Review of Economic Studies*, 57(2): 279–298.

Jäger, K. 2016. "EU KLEMS Growth and Productivity Accounts 2016 release - Description of Methodology and General Notes."

Karabarbounis, L., and B. Neiman. 2014. "The Global Decline of the Labor Share." *The Quarterly Journal of Economics*, 129(1): 61–103.

Kreps, D. M. 1998. "Anticipated utility and dynamic choice." *Econometric Society Monographs*, 29: 242–274.

Krueger, A. B. 1999. "Measuring labor's share." *American Economic Review*, 89(2): 45–51.

Kurman, A., and E. McEntarfer. 2017. "Downward Wage Rigidity in the United States: New Evidence from Administrative Data." mimeo.

Kurozumi, T., and W. Van Zandweghe. 2012. "Learning about monetary policy rules when labor market search and matching frictions matter." *Journal of Economic Dynamics and Control*, 36(4): 523–535.

Mankiw, N. G., R. Reis, and J. Wolfers. 2004. "Disagreement about inflation expectations." In *NBER Macroeconomics Annual 2003, Volume 18*. The MIT Press, Cambridge, MA, 209–270.

Marcet, A., and T. J. Sargent. 1989. "Convergence of least squares learning mechanisms in self-referential linear stochastic models." *Journal of Economic Theory*, 48(2): 337–368.

Mareek, P., and E. Orgiazzi. 2015. "Development and the labor share." THEMA Working Paper 2015-09, Universite de Cergy-Pontoise, Working Paper.

Martins, P. S., G. Solon, and J. P. Thomas. 2012. "Measuring What Employers Do about Entry Wages over the Business Cycle: A New Approach." *American Economic Journal: Macroeconomics*, 4(4): 36–55.

McCallum, B. T. 1983. "On non-uniqueness in rational expectations models: An attempt at perspective." *Journal of Monetary Economics*, 11(2): 139–168.

Merz, M. 1995. "Search in the labor market and the real business cycle." *Journal of Monetary Economics*, 36(2): 269–300.

Milani, F. 2007. "Expectations, learning and macroeconomic persistence." *Journal of Monetary Economics*, 54(7): 2065–2082.

Milani, F. 2011. "Expectation Shocks and Learning as Drivers of the Business Cycle." *Economic Journal*, 121(552): 379–401.

Miller, R. E., and P. D. Blair. 2009. *Input-Output Analysis*. Cambridge University Press.

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

Mortensen, D. T., and C. A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies*, 61(3): 397–415.

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

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

Pfajfar, D., and E. Santoro. 2010. "Heterogeneity, learning and information stickiness in inflation expectations." *Journal of Economic Behavior & Organization*, 75(3): 426–444.

Pissarides, C. A. 2000. *Equilibrium unemployment theory*. The MIT Press, Cambridge, MA.

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

Pissarides, C. A., and B. Petrongolo. 2001. "Looking into the Black Box: A Survey of the Matching Function." *Journal of Economic Literature*, 39(2): 390–431.

Psacharopoulos, G. 1994. "Returns to investment in education: A global update." *World Development*, 22(9): 1325–1343.

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

Rogerson, R., and R. Shimer. 2011. "Search in Macroeconomic Models of the Labor Market." *Handbook of Labor Economics, Volume 4A:* 619–700.

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

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

Shimer, R. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review*, 95(1): 25–49.

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

Smith, J. C. 2011. "The Ins and Outs of UK Unemployment." *Economic Journal*, 121(552): 402–444.

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

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

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

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

Timmer, M. P., E. Dietzenbacher, B. Los, R. Stehrer, and G. J. Vries. 2015. "An Illustrated User Guide to the World Input-Output Database: The Case of Global Automotive Production." *Review of International Economics*, 23(3): 575–605.

Valentinyi, A., and B. Herrendorf. 2008. "Measuring factor income shares at the sectoral level." *Review of Economic Dynamics*, 11(4): 820–835.

Voskoboynikov, I. B. 2012. "New measures of output, labour and capital in industries of the Russian economy." GGDC Research Memorandum 123, Groningen Growth and Development Centre.

World Bank. 2008. *Global purchasing power parities and real expenditures: 2005 international comparison program*. World Bank Group.

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