Flexibility of New Hires' Earnings

Reamonn Lydon*[1] and Matija Lozej†[1]

[1]Central Bank of Ireland

Abstract

Two recent papers that distinguish between the wage flexibility of new hires, incumbents and job changers in the U.S. yield opposite results. Combining Household Finance and Consumption Survey and tax data for Ireland, we find that earnings of new hires from unemployment are substantially more flexible compared to earnings of incumbent workers or job changers. The findings are robust. Earnings of new hires from unemployment are more procyclical for workers with less valuable outside options. Our results indicate that wage rigidity may not be a suitable device to generate unemployment volatility in macroeconomic models.

JEL classification: C23, E24, J31, J52.

Keywords: Wage bargaining, Wage rigidity, Employment fluctuations, Business cycles.

---

*Reamonn.Lydon@centralbank.ie

†Matija.Lozej@centralbank.ie. The views expressed in this paper are the authors' only and do not represent the views of the Central Bank of Ireland or ESCB. We thank participants at the 2016 Irish Economic Association Conference, Wouter den Haan, Luca Onorante, Gabriel Fagan and the referee for helpful comments. We would particularly like to thank Paul Crowley and Gerard Reilly at the Central Statistics Office, Ireland for providing access to the data. The CSO does not take any responsibility for the views expressed or outputs generated from this research.

1 Introduction

Since Shimer (2005) pointed out that a standard Diamond-Mortensen-Pissarides search and matching model cannot replicate the volatility of labour market variables observed in the data, several ways to address this have been proposed. One of the remedies, suggested already in Shimer's 2005 paper, is incorporating wage rigidity into the model setup. But while aggregate wages seem to be sticky (Bewley, 1999), when it comes to the firm's hiring decision, the wage flexibility that matters is that of the new or marginal worker (Pissarides (2009)), and not of the aggregate wages. The new hires' wage rigidity is thus an important amplification mechanism in a whole class of models, which has spurred a renewed interest in this issue in the recent literature.

Starting with Bils (1985) there is an empirical literature that tries to identify the cyclical sensitivity of the wages of new hires separate from existing workers. Much of this literature fails to distinguish between the wages of new hires who were previously unemployed and the wages of new hires who were previously employed ('job changers'). The main reason is the lack of data to allow researchers to clearly identify the two types of new hires. Two notable exceptions are the papers by Haefke et al. (2013) (HSvR) and Gertler et al. (2015) (GHT) on wages of new hires from unemployment in the US. Interestingly, both papers reach opposite conclusions: HSvR find that new hires' are more sensitive to the business cycle (productivity shocks); GHT, on the other hand, find no evidence of greater wage flexibility for new hires from unemployment.

In this paper, we investigate the sensitivity of the weekly earnings of new hires in Ireland to changes in local unemployment rates. Crucially, we distinguish whether new hires were from unemployment or inactivity (i.e. truly 'new hires') or from other jobs ('job changers'). To do so, we exploit a new administrative panel dataset on the pre-tax earnings of employees in Ireland over the period from 2005 to 2014, linked to the 2013 Household Finance and Consumption Survey. Our main findings are that new hires' weekly earnings are substantially more procyclical than the earnings of incumbent workers and that this procyclicality is driven by workers transitioning from from unemployment or inactivity. For job changers, weekly earnings do not appear to exhibit different cyclical behaviour than for incumbent workers.

Mainly to address concerns around composition bias affecting our results, we subject the data to a wide range of robustness checks – including an examination of the cyclical behaviour of the pay of new hires and incumbents *within* age, education, industry and occupation groups. Across all workers, our estimated elasticity of new hire pay to unemployment ranges from -1.0 to -1.7; that is, a 1% increase in the unemployment rate reduces new hire pay by 1 to 1.7%. These results are consistent with the findings of HSvR for the US – who appeal to long-term contracting and insider/outsider theories of wage determination to explain their results – but are in stark contrast with the results in GST. We also find that earnings of new hires from unemployment are more flexible when workers have outside options that are less valuable, for instance when they are less educated or older, but not old enough to be able to wait out the unemployment spell until retirement.

Apart from providing evidence on the flexibility of new hires from a country other than the US, the investigation of Irish data is interesting because our dataset spans an extraordinary period in Irish economic history (see Figure 2). The first two years of the sample catch the end of a prolonged period of very low unemployment (averaging 4.5%) which coincided with the Celtic Tiger and credit boom years. This contrasts with earlier decades which were characterised by persistently high unemployment and non existent employment growth (O'Connell, 1999). With the onset of the recession in 2007, which had turned into a full-blown financial crisis by 2010, unemployment rose rapidly, peaking at just under 15% in 2012. As Table 1 in Section 2 shows, more than half of the 300,000 jobs lost during the recession were in construction – highlighting the extent to which the economy was reliant on this one sector during the period of the credit boom in particular (from 2003 onwards). On the back of an exceptionally strong export performance, employment growth resumed in 2013, such that by end-2014 employment was 76,000 above trough levels. Our earnings data also covers this remarkable recovery phase. We discuss wage and earnings developments in both the aggregate and our micro data in more detail in the data section below.

To a certain extent, the depth of the recession in Ireland provides an ideal setting in which to examine the issue of cyclical wage or earnings flexibility. There is some evidence that during stronger recessions wages in general may lose some of their downward nominal rigidity; see for example Abbritti and Fahr (2013); Fagan and Messina (2009) and Fabiani et al. (2015). In the case of Ireland during the period in question, we argue that changes to the institutional wage setting architecture, and in particular the abandonment of collective bargaining agreements which had been in place since the late 1980s and which in part protected new hires earnings, may have played a key role. This variation in the data should help us to better identify the cyclicality of wages. For instance, the result in Haefke et al. (2013) that wages of new hires from unemployment are procyclical is barely statistically significant, which may be due to their dataset that mostly spans the Great Moderation period. This may also be the reason why the results of GHT, whose sample includes the Great Recession, differ from the results of HSvR. On the other hand, one of the downsides of looking at labour markets during a deep recession is that in countries that experience large negative labour demand shocks, the wages of *existing* workers might also become more flexible, so the difference between rigidity of wages of existing and newly hired workers might be more difficult to identify.

The remainder of the paper is organised as follows: in Section 2 we describe the data, summarise trends in earnings and compare labour market transition rates in our earnings database with those in the Labour Force Survey (LFS).[1] Section 3 presents our empirical framework for testing the cyclicality of new hires' earnings depending on the their previous work status. Section 4 concludes.

[1]The Irish Labour Force Survey was previously called the *Quarterly National Household Survey*. The survey is a five-quarter rolling panel of approximately 20,000 individuals per quarter. Currently, the LFS does not ask individuals about their level of earnings.

2 Data & trends

2.1 Earnings database

The earnings data we use is taken from an administrative tax database on the annual earnings of employees in Ireland from 2005 through to 2014. This panel data also contains information on weeks worked each year, thereby allowing us to estimate average weekly earnings. There are two key advantages to the data we use for the empirical analysis. Firstly, it is data on earnings drawn from individual tax records, and thus should be largely free of any measurement error issues which may bias results. Secondly, it is a panel dataset with up to 10-years of individual earnings data, meaning that we are able to control for individual fixed effects which could bias the results. For example, if firms' hiring standards are also pro-cyclical such that the composition of newly-hired workers changes through the cycle, as suggested in Sedláček (2014), then it will be important to control for this in the regression.

The earnings database contains no information on worker or employer characteristics. In order to get this data, which is important in order to be able to control for individual heterogeneity (see Haefke et al. (2013)), we link the individuals in the tax database to the same individuals in a household survey carried out in 2013 – the *Household Finance and Consumption Survey* or HFCS. The HFCS was carried out by the Irish *Central Statistics Office* from March to September 2013, mainly for the purposes of collecting data on household wealth; see Lawless et al. (2015) for a description of the survey.[2] However, it also contains extensive information on individual characteristics – including education, experience, job history and job tenure – as well as information about the job itself, such as hours (in 2013), occupation and sector. We call the linked HFCS-tax records dataset the 'HFCS-Admin' data.

[2]The matching of ten-years of earnings from administrative data to individuals and households in the HFCS was carried out by the CSO using personal identifiers common to both datasets.

The tax database lists the social insurance category to which the worker belongs in any given year. In Ireland, workers can pay different social insurance contributions depending on their sector of work, level of earnings, retirement status and previous employment (e.g. military, etc). Given our focus on the wage-setting behaviour of firms we restrict our attention to workers in industrial, commercial and service-type employment with gross earnings of €38 or more a week from work.[3]

Table 1, which gives an overview of the data, shows that in any given year we observe the earnings of between 4,500 to 4,800 workers. During the recession, some *new* public sector workers such as teachers and medical professionals in training had their pay cut relative to incumbents, resulting in what has been termed 'two-tier' pay scales.[4] To see whether our results are driven by public sector pay cuts, we run our analysis both with and without public sector workers. The final column in Table 1 shows that excluding public sector workers leaves a sample of around 4,000 workers each year.

The HFCS is a household survey designed to be representative of population characteristics, earnings and wealth in 2013. However, comparisons with LFS data shows that the HFCS-Admin dataset does an excellent job of tracking labour market over the ten years of data. For example, developments in the percentage of employees in the LFS and HFCS-Admin (i.e., earners in the administrative data) are very similar (Figure 1). This reassures us that the HFCS-Admin data captures well the the cyclical shifts in the labour market over period in question.

[3]These are social insurance category 'A' workers, which covers almost all private sector workers, as well as most public sector workers hired after 1995. Of the 44,541 employee-year earnings observations in the data, 42,136 (95%) are social insurance category A workers. The largest omitted grouping are category J workers – workers earning less than €38 per week, or workers on State-sponsered job training schemes. See The Department of Social Welfare webpage for further background on social insurance classes in Ireland.

[4]For example, new teachers starting pay was cut by almost 20% from 2010 onwards. See The Irish Times, Tuesday 1 March 2016 for an overview.

Table 1. Earnings observations in the HFCS-Admin dataset

	Male	Female	Total (All Class 'A')	Total (Ex-public sector)
2005	2,267	2,166	4,433	3,758
2006	2,320	2,276	4,596	3,906
2007	2,461	2,343	4,804	4,070
2008	2,487	2,370	4,857	4,112
2009	2,393	2,246	4,639	3,905
2010	2,366	2,174	4,540	3,796
2011	2,385	2,140	4,525	3,771
2012	2,427	2,217	4,644	3,875
2013	2,478	2,271	4,749	3,979
2014	2,454	2,328	4,782	4,027
2006-2014	21,771	20,365	42,136	35,441

Figure 1. Percentage of employees in the sample, by year

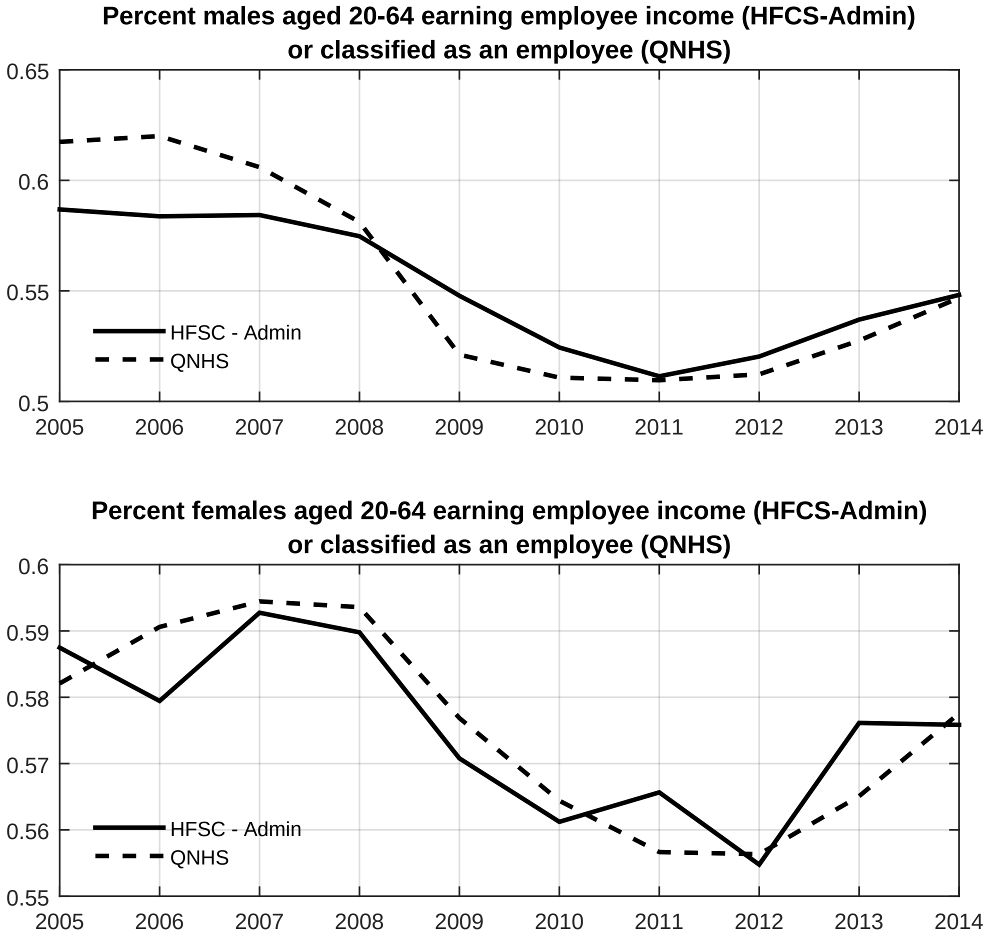

Notes: The LFS (QNHS) variable we use to identify employees is STAPRO. The values are year averages across quarters using sample weights. Because the HFCS-Admin is a representative household survey in 2013, we restrict the sample to Irish born adults to try and account for the propensity for higher emigration rates by previous immigrants.

We use the labour market history information in the HFCS along with the administrative earnings data to identify gaps in employment histories and categorise workers in any given year into three categories:

1. Workers who transition from inactivity or unemployment to employment from one year to the next ('New Hires');
2. Workers who change job, but it is an employment-to-employment transition, with no (extended) spell of unemployment or inactivity in between ('Job changers'); and
3. Workers who do not change job from one year to the next.

Using some examples from the data, Figure 11 ('New Hires') and Figure 12 in the appendix illustrate the methodology we use to identify workers in each group. New hires from unemployment or inactivity are identified primarily from the administrative earnings data. As Figure 11 shows, a worker who exhibits zero earnings in the previous year is assumed to have made the transition from unemployment or inactivity to employment.

'Job changers' are identified from both the HFCS survey information *and* the administrative earnings data. In the HFCS, workers are asked to state the year and month they started their current job. If a worker had earnings in the year *before* they started their job, we categorise them as job changers, if not, they are categorised as new hires from unemployment. Job changers experience a significant increase in real earnings upon starting a new job, around 6% across all years. These changes are, not surprisingly, also procyclical (See Figure 13 at the end of the paper).

2.2 Aggregate cyclicality of earnings of different groups

The data described above allow us to construct a time series for the aggregate earnings of new hires (either from job-to-job or from unemployment), and we use these data to analyse their cyclical behaviour compared to earnings of those who did not change jobs. The ten-year period spans the end of the previous credit-fuelled expansion, the contraction/stagnation years from 2008 to 2012, and the recovery phase from 2013 onwards. Figure 2 and Table 1 below paint a picture of labour market trends in the run-up to the crash, during the recession and in the more recent recovery phase.[5]

An important institutional development during this period that is likely to have influenced wage behaviour was the move away from a centralised wage bargaining process – the collective bargaining arrangements known as *Partnership Agreements*, which have been in place since the late 1980s. These agreements were reneged upon when the crisis hit in 2008.[6] For instance, the pay increases scheduled under the second phase of the *Towards 2016* national wage agreement were postponed indefinitely for the public sector in 2009. A formal agreement on pay determination to replace the *Towards 2016* agreement, which formally expired during 2010, was not put in place and therefore the social partners have been operating without a formal agreement on pay determination since 2010. Importantly, these agreements stipulated that new hires should receive the same terms as existing (similar) employees, including wages. Abandonment of these agreements may, in part, have may have made it easier for firms to adjust the wages of new hires during the downturn. Note that these developments should not be viewed as exogenous institutional changes, but rather as endogenous developments resulting from the severity of the downturn.

There is also international evidence on changes in the sensitivity of wages of new hires, which is available in the most recent *Wage Dynamics Survey* (2014) conducted by Central Banks participating in the ECB's *Wage Dynamics Network*. In this cross-country survey of firms, managers were asked whether during the 2010-13 period it became "easier to lower the wages at which new employees were hired". Figure 3 shows that firms in countries that experienced a larger increase in unemployment during 2010-13, such as

[5]An online appendix to this paper provides a more comprehensive set of business cycle summary statistics on the key variables of interest.

[6]Leddin (2010) provides a historical overview of the agreements.

Ireland, Greece, Spain and Cyprus, are significantly more likely to say that it became *easier* to lower the wages at which new employees were hired.[7]

[7]See the website of the Wage Dynamics research Network for background to the survey. Linehan et al. (2015) provide specific details for Ireland, including a copy of the common cross-country questionnaire.

FIGURE 2. Employment, hours, unemployment and GDP trends in Ireland

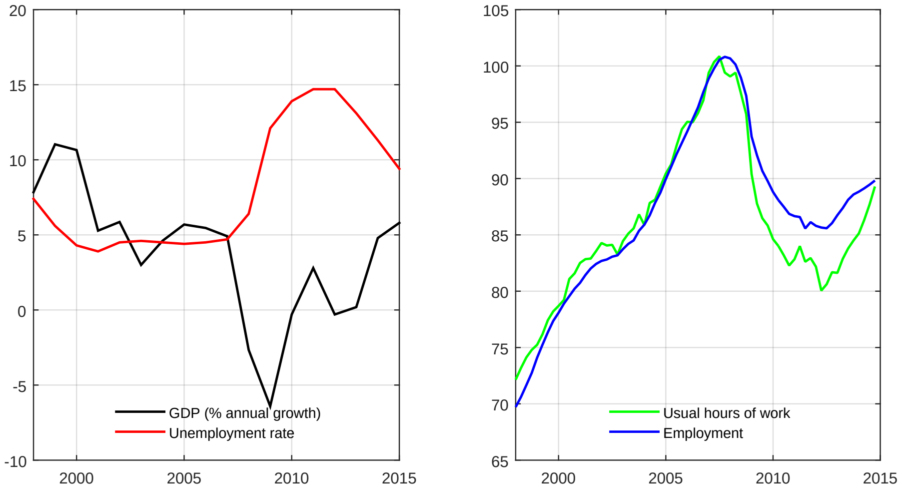

Source: Central Statistics Office. GDP data for 2015 is the Central Bank Q4 Bulletin forecast. Unemployment data for 2015 is to September 2015

Table 2. Employment by sector ('000s)

	2003	2007	2012	2014
All NACE economic sectors	1,810.1	2,143.1	1,837.9	1,913.9
Industry (B to E)	300.4	299.1	234.0	239.0
Construction (F)	181.8	270.3	101.8	109.4
Services (G to U)	1,205.3	1,457.3	1,413.9	1,453.5

Source: Irish Labour Force Survey (QNHS). Figures are within year means.

Using Labour Force Survey data, Figure 4 shows that the proportion of workers who start a new job in a given period tends to track GDP growth (positively) and unemployment (negatively). The correlation with GDP growth is 0.88 (p-value = 0.000) and with unemployment -0.80 (p-value = 0.001).

As we are particularly interested in whether a worker who starts a new job comes from unemployment/inactivity ('new hire') or a previous job ('job changer'). In Figure 5 we separate out the two groups of new workers, and focus on just the new hires. The chart shows that new hire rates in the LFS closely track those in the constructed HFCS-Admin data. The new hire rate also broadly tracks GDP trends, with a correlation coefficient of 0.80 (p-value=0.001) between GDP growth and the new hire rate in the HFCS-Admin dataset.

Figure 3. Evidence on the procyclical wage flexibility of new workers in the Wage Dynamics Survey

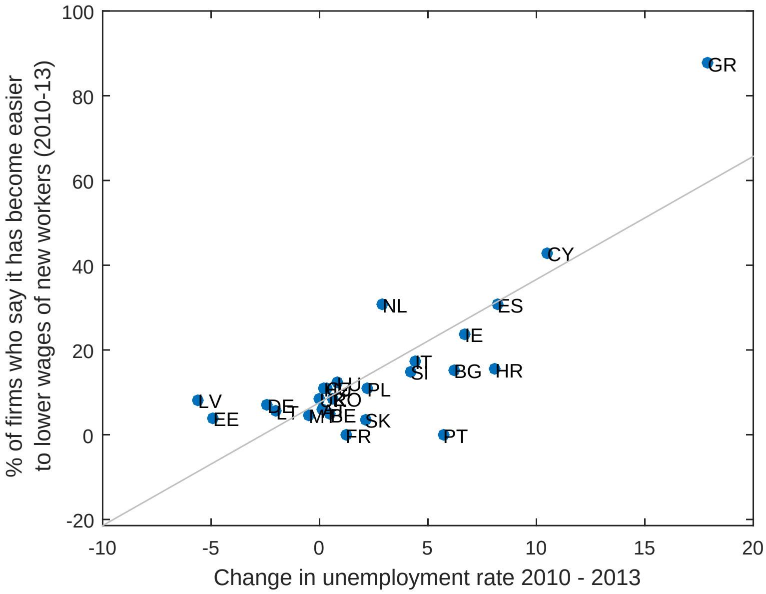

*Source:*Wage Dynamics Network data, question C3.4.

Figure 4. Cyclicality of job starts

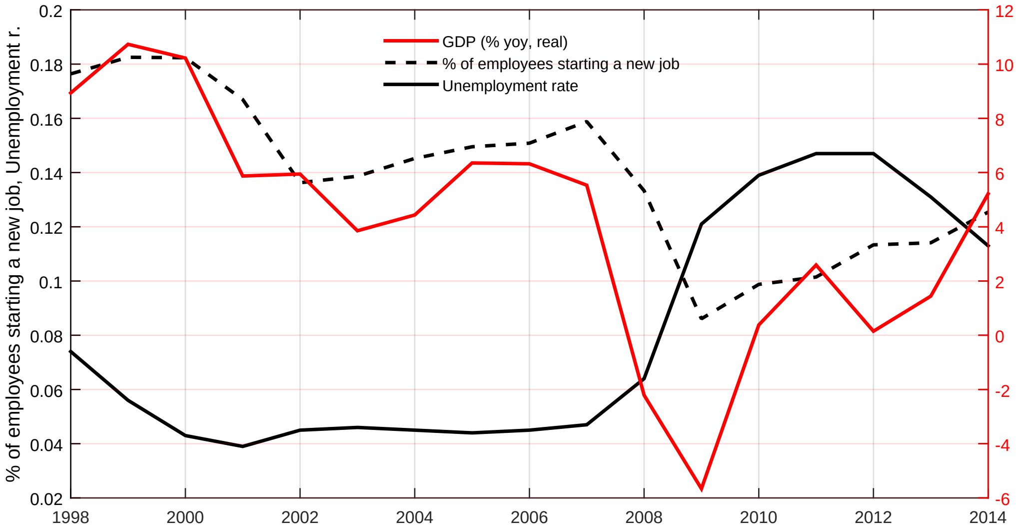

Source: LFS. To avoid issues around seasonality, we use the third quarter in each year only.

Figure 5. Transitions to employee status from unemployment/inactivity or self-employment

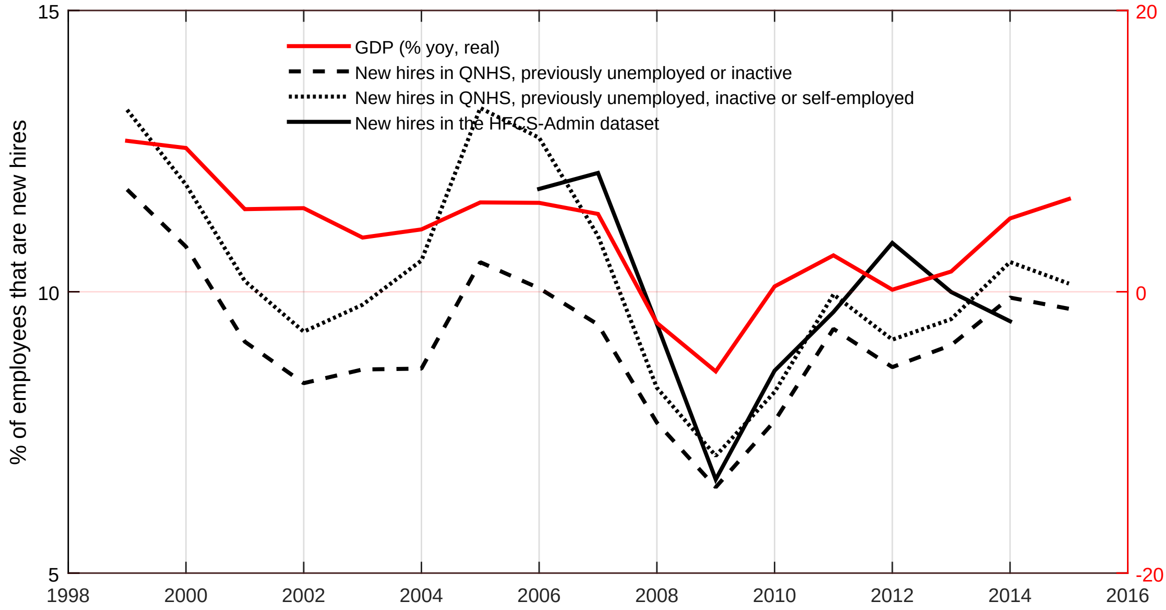

Source: LFS (formally, the 'QNHS') and HFCS-Admin. In the LFS we use the four quarter lagged values for STAPRO to identify transitions. In the HFCS-Admin data, a transition is simply an observation from not-earning to earning.

Income from self-employment is not recorded in HFCS-Admin, therefore some of those we identify as transitioning from ‘not earning’ to earning may have been self employed prior to becoming employees. To determine whether this missclassification is an issue for our results, Figure 5 uses the LFS to separate the new hire rate into those coming from unemployment/inactivity and those coming from self-employment (these are one-year transition rates). For our purposes, it is noteworthy that (a) the share of self-employed to employee transitions is very small – on average 10% of new hire employee transitions are from self-employment; and (b) whilst the overall pattern does display some cyclicality, the share of previously self-employed employees does not appear to be strongly pro-cyclical, particularly in the the down-turn from 2007-10. This is borne out by the correlation coefficient: the correlation between the new hire rate from self-employment to employee and GDP growth is just 0.47 (p-value 0.060). We argue that whatever misclassification exists in HFCS-Admin, it is likely to be small and any potential bias problems are also small.[8]

Our database does not provide annual information on hours worked. This means that the dependent variable in our empirical analysis is weekly earnings. To the extent that the cyclical behaviour of hours worked may differ by new hire type (and also versus existing workers) this might bias our results. Figure 6, from the LFS, shows that the hours of *all* worker-types are pro-cyclical.[9] However, the data does not suggest that the hours of new hires are significantly more or less pro-cyclical than other groups – particularly during the 2007-10 period which is the period where new hires’ earnings fall the most.

[8]It is difficult to know exactly what the direction of the bias might be in this case. If previously self-employed employees are more likely to have higher earnings than previously inactive or unemployed employees, then there will a downward bias to our estimates of the cyclicality of new hires earnings, or vice-versa if the opposite is true.

[9]The LFS data shows that the proportion of workers in part-time employment (working up to three days per week) rose from 15% to 22% of those employment between 2007 and 2012.

Figure 6. Actual hours worked by existing workers and new hires (2007=100)

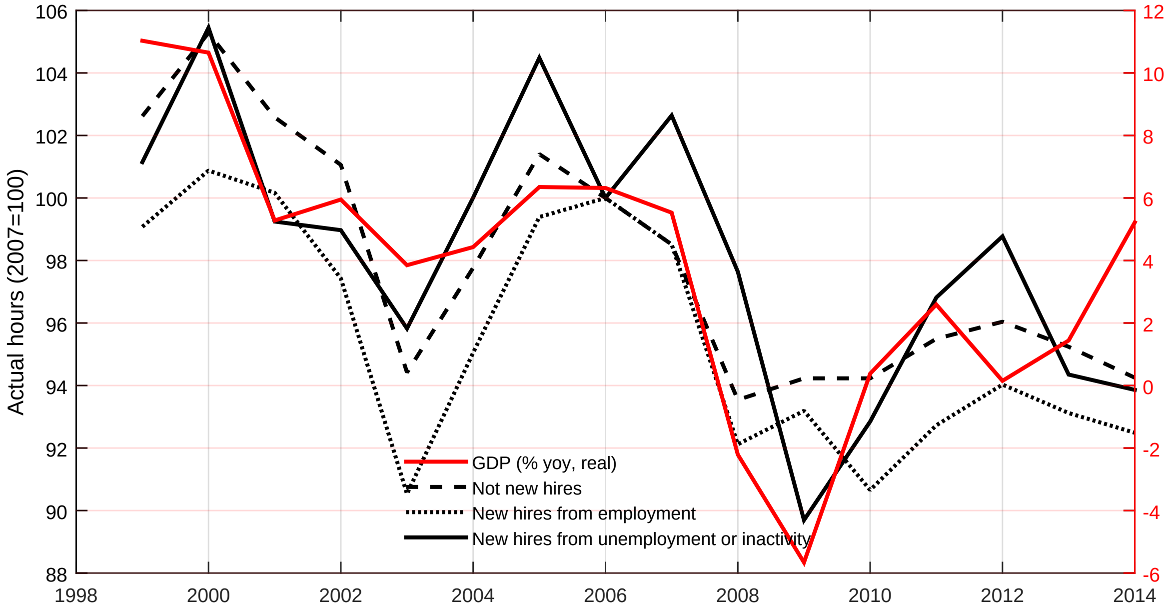

Source: LFS (formally, the 'QNHS').

For further assurance, we use HFCS-Admin to track weekly and hourly earnings at the aggregate level for new hires versus existing workers. The sample for this comparison is all employees, as shown in the third column of Table 1. Figure 7 shows that there is a clear pattern of declining earnings for new hires throughout the recession, when compared with existing workers or job changers. At the low-point of the sample, in 2013, new hires' pay had declined by over 20% from 2007 levels, with the pay of other groups reduced by around 5%. Table 8 at the end of the paper, which shows the results from a difference-in-difference earnings regression, shows that these differences are statistically significant. Figure 8 shows a similar chart, but this time for hourly earnings, using mean actual hours worked for each group from LFS. We emphasise that the dynamics of hourly earnings for all groups remains unaffected, which gives us more assurance that variations in hours worked between groups did not play an important role in the sensitivity of weekly earnings of new hires. These graphical comparisons make no allowance for compositional changes or other factors such as skills (see Haefke et al. (2013)). In the next section we estimate several regressions to both quantify this difference and see whether it remains after controlling for worker heterogeneity.

FIGURE 7. Weekly earnings of new hires versus existing workers

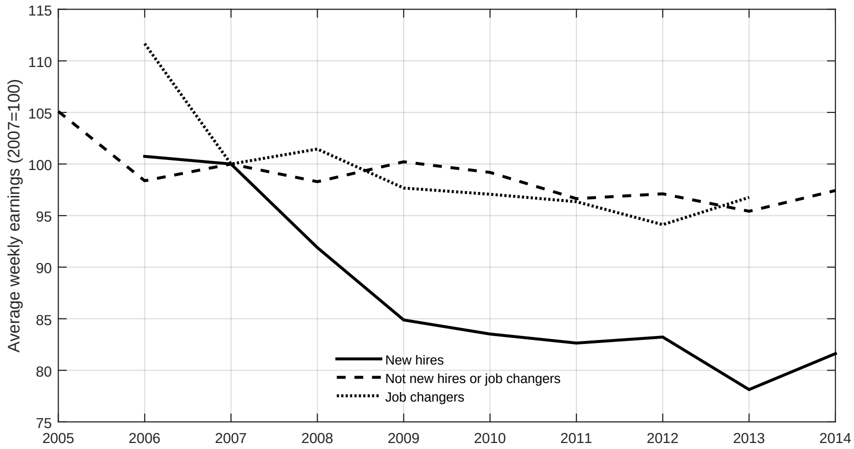

Source: HFCS-Admin. Weekly earnings deflated by the consumer price index. Earnings of workers aged 20-64, working a minimum of 26 weeks in the year.

FIGURE 8. Hourly earnings of new hires versus existing workers

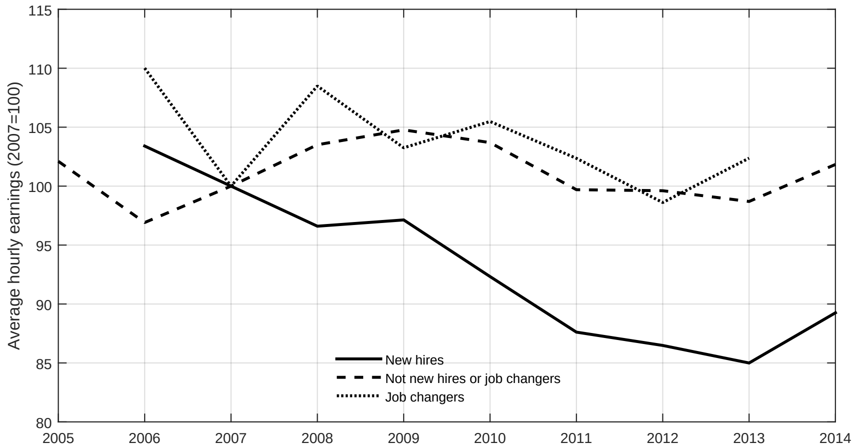

Source: HFCS-Admin, LFS. For each group: Average weekly earnings deflated by the consumer price index, divided by mean actual hours from LFS. Earnings of workers aged 20-64, working a minimum of 26 weeks in the year.

FIGURE 9. Co-movement of unemployment and the ratio of existing workers' earnings to new hires' earnings

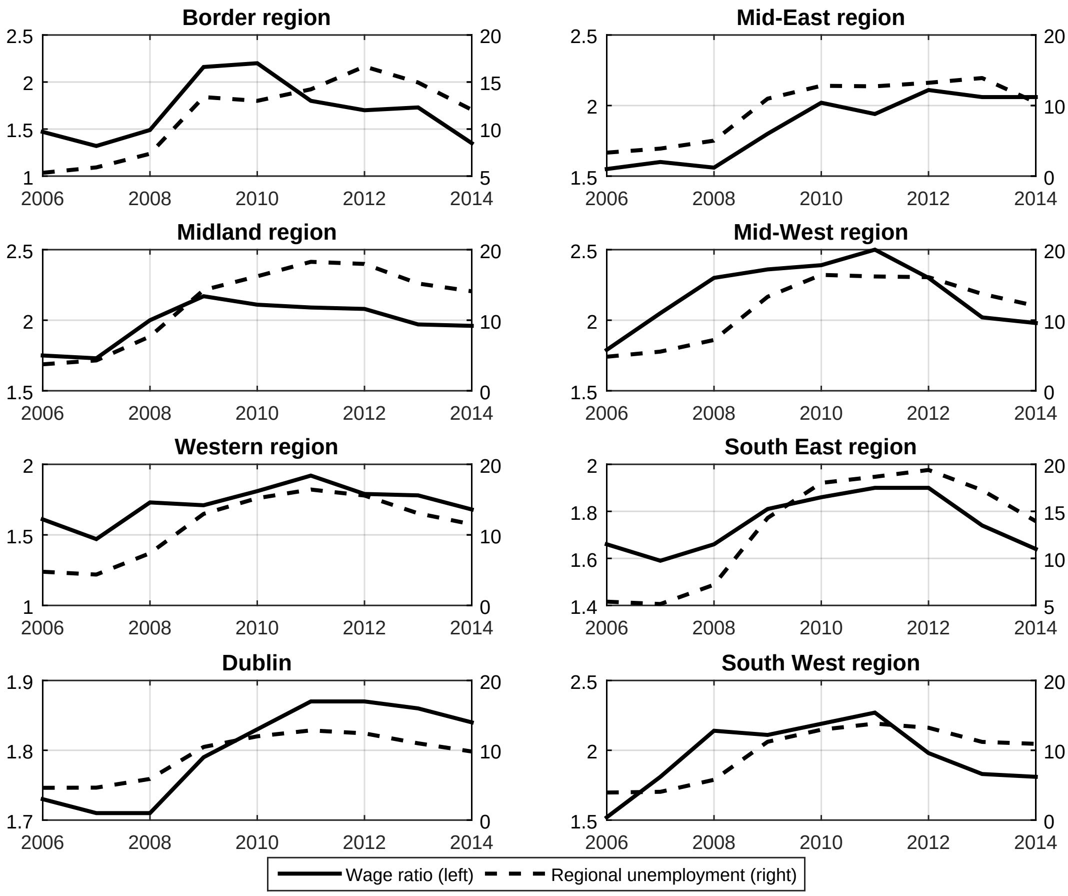

Source: HFCS-Admin & LFS for regional unemployment rates.

Finally, we look at the co-movement of earnings and unemployment over time. This sets-up the regressions in the next section, where the regional unemployment rate is the variable used to capture changes in the business cycle. Figure 9 compares the evolution of the ratio of existing workers' weekly earnings to new hires' weekly earnings with the unemployment rate for eight regions in the data.[10] For some regions, there is a strong correlation between the ratio of the earnings of the two groups and the unemployment rate, e.g. Dublin, Mid-East, South-East, South-West and Midland (correlation coefficient of 0.85 or higher); whereas for others the relationship appears to be weaker, e.g. Border and the Mid-West – although in the latter two cases, there is still a significant positive correlation (0.54 and 0.66 respectively). The chart also motivates our using regional as opposed to national unemployment rates in the regressions. At the lowest point of the cycle, in 2012, there was a considerable difference in the unemployment rates across regions. For example, Dublin, where the rate was lowest, saw unemployment averaging 12.4%, whereas in the South-East in the same year, it was pushing 20%.

3 Econometric specification and estimation results

Our regression specification closely follows Bils (1985) and the baseline specification used more recently by GHT (2015). The measurement equation for the log weekly earnings of individual i in region j at time t is as follows:

$$w_{ijt} = x'_{ijt}\beta_x + \beta_u u_{jt} + \beta_{new} \cdot new_{ijt} + \beta_{new,u} \cdot new_{ijt} \cdot u_{jt} + \alpha_i + \eta_{ijt}, \tag{1}$$

This specification is also closely related to the empirical literature on wage equation estimation using micro data; see for example, Topel and Ward (1992), Barlevy (2001), Barlevy (2002), Solon et al. (1994) and Martins et al. (2012). In the regression u_{jt} is the

[10]For this chart, we group job changers with existing workers. The patterns are unchanged if this group is excluded, as Figure 7 suggests. The eight regions are the NUTS III regions for Ireland: Border, Midland, Western, Dublin, Mid-East, Mid-West, South East and South West.

regional unemployment rate, new_{ijt} equals one if the worker has recently (i.e. in the last year) started the job and zero otherwise. At this stage, we do not differentiate between new workers that are new hires from unemployment/inactivity and job changers. The variables α_i and η_{ijt} are individual fixed effects and a random error term respectively. The unemployment variable, which is year (2005-2014) and region ($\times 8$, see footnote 10) specific, captures the influence of the business cycle on wage setting. Recall that the country-wide collective agreement was abandoned at the onset of the crisis. With regional unemployment rates we can better capture local cyclical labour market conditions in an environment where labour is not perfectly mobile across regions. Devereux (2001) contains a discussion of the pros and cons of including country versus regional unemployment rates. He opts for the former, arguing that measurement error in regional unemployment rates might introduce a negative bias. We also estimate specifications using national unemployment rates (results available on request) and obtain very similar overall results. The interaction with the new hires dummy variable is included to capture the possibility that the earnings of new hires are more sensitive to the business cycle.

The key feature of our data is that it enables us to differentiate between the earnings of new hires from (previous) employment, that is 'job changers', captured by the dummy variable $newE_{ijt}$, and new hires from unemployment ($newU_{ijt}$). To capture the differences, the baseline specification in equation (1) is augmented with additional interaction terms and follows the specification of equation 2 in GHT (2015):

$$\begin{aligned} w_{ijt} = x'_{ijt}\beta_x + \beta_u u_{jt} & \\ & + \beta_{newU} \cdot newU_{ijt} + \beta_{newU,u} \cdot newU_{ijt} \cdot u_{jt} \\ & + \beta_{newE} \cdot newE_{ijt} + \beta_{newE,u} \cdot newE_{ijt} \cdot u_{jt} \\ & + \alpha_i + \eta_{ijt}. \end{aligned} \tag{2}$$

If the earnings of new hires are *more* pro-cyclical than earnings for other workers, then β_{new} and β_u in equation 1 are both negative and statistically significant. Moreover, if earnings of new hires from unemployment or of job changers behave differently in response to the cycle, then this would be captured by $\beta_{newU,u}$ and $\beta_{newE,u}$.

The first column in Table 3 shows the results from estimating equation (1) as a fixed effects regression. The negative coefficients on the unemployment rate (β_u) and the new hire-unemployment rate interaction suggests that the earnings of new workers are significantly more sensitive to the economic cycle. Column (2) separates out new workers into those that were previously unemployed or inactive (true 'new hires' in our taxonomy) and new workers who transition from another job. The results suggest that it is the heightened sensitivity of new hires' earnings to the unemployment rate which drives the earlier result. We find that the semi-elasticity of new hires' weekly earnings with respect to the unemployment rate is -1.554 ($\beta_u + \beta_{newU,u}$) , whereas for all other workers (incumbents and job-to-job changes) the semi-elasticity is -0.376. In other words, for new hires a 1% increase in the unemployment rate is associated with a reduction in weekly earnings of 1.554%. In the final column we drop all public sector workers, for the reasons discussed earlier. As might be expected, in general we observe a greater sensitivity of private sector workers' earnings levels to the unemployment rate, with a semi-elasticity of -0.464, compared with the earlier estimate of -0.376. However, the interaction with new hires is more-or-less unchanged, at -1.206. Overall, the semi-elasticity excluding public sector workers (-0.464-1.206=-1.67) is marginally higher than the specification including these workers, although the differences are not statistically significant.

Our results regarding the sensitivity of earnings of new hires from unemployment to the business cycle are in stark contrast with the results for the U.S. in GHT (2015), who find that wages of new hires from unemployment are not more sensitive to the cycle than wages of stayers when one controls for job changers. Moreover, they find that wages of job changers are more sensitive to the cycle and go on to control for the composition effect, after which the excess sensitivity of wages of job changers to the business cycle disappears. We do not find such effects in the first place. Our results indicate that the opposite is the case in Ireland - it is the stronger sensitivity of earnings of new hires from unemployment to the business cycle that drives the sensitivity of earnings of new hires. Our results are in line with the results in HSvR (2013), who find $\beta_{newU,u} < \beta_{newE,u} < 0$, albeit only marginally statistically significant. Unlike their results, ours are statistically significant and robust. In the remainder of the paper we focus on examining the robustness of this result and on providing more details on where it comes from.

Table 3. Unemployment and the earnings of new hires

Log(weekly earnings)	(1)	(2)	(3)
	Fixed effects	Fixed effects	Fixed effects (ex-public sector)
Unemployment (β_u)	-0.376***	-0.376***	-0.464***
	(0.0688)	(0.0687)	(0.0791)
Newhire #Unemp ($\beta_{new,u}$)	-0.703***		
	(0.130)		
Newhire, U #Unemp ($\beta_{newU,u}$)		-1.178***	-1.282***
		(0.156)	(0.1696)
Newhire, E #Unemp ($\beta_{newE,u}$)		0.0421	0.0189
		(0.211)	(0.2416)
Age	0.0945***	0.0939**	0.0960***
	(0.00259)	(0.00259)	(0.0029)
Age2	-0.00105***	-0.00105***	-0.00127***
	(0.00003)	(0.00003)	(0.00004)
Constant	4.123***	4.145***	4.173***
	(0.0531)	(0.0532)	(0.0592)
Semi-elasticity (dw/du)	1.08	1.55	1.75
Observations	42,136	42,136	35,441
Number of id	6,775	6,775	5,968
R-squared	0.188	0.190	0.22

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

3.1 Alternative specifications and compositional shifts

Our baseline fixed-effects regression suffers from two potential mis-specification problems widely discussed in the literature. The first is a spurious regression problem which could arise if earnings and the unemployment rate share a common time trend and are therefore integrated (see, for example, the discussion in HSvR and Solon et al. (1994)).

The second potential issue is that our results may be due to a shift in the composition of new hires during a recession. For example, if there are fewer highly skilled or highly productive workers amongst new hires during a recession, then this might reduce the earnings of new hires on average. Solon et al. (1994) call this a countercyclical omitted variable bias. In the case of Ireland, where tens of thousands of young graduates and perhaps higher-skilled workers emigrated during the recession, this is obviously a concern.[11] These compositional shifts are evident in the HFCS-Admin data. For example, from 2006-2008, 42% of new hires were college graduates; by 2012, this had fallen to 36%. There are similar shifts in the average age of new hires: pre-2008, the average age of a new hire was 35.5, by 2012 this had fallen to 33.

A first difference specification can be used to address both issues. However, as HSvR point out, the downside of doing this with new hires is that we lose the first observation, which is, naturally, of great interest in this case. Another option is to re-estimate the fixed-effects specification within homogeneous sub-groups, such as age, education, occupation or sector groups. One issue with this approach is that it only captures differences in observables. The following sections present estimates of the earnings-unemployment semi-elasticity using both approaches.

[11]See Glynn et al. (2013) for an overview of the education and previous work experience of Irish emigrants during the recession. The authors note that, relative to emigrants' profiles in earlier decades, "...today's emigrants are much more likely to have a high standard of education than the population in general and arguments referencing a 'brain drain' are not misplaced." (page 29).

3.1.1 Estimation in first differences

Both Devereux (2001) and Solon et al. (1994) estimate a micro counterpart to macro wage regressions in first differences. As noted above, a difficulty which arises when we look at new hires is that there is no $t-1$ observation of earnings for that particular sub-group. Haefke et al. (2013) propose a two-stage solution: in the first stage, generate a wage index for group g (i.e. new hires or incumbents) which relates the average wage of group g workers, w_{gt}, as follows:

$$log\, \hat{w}_{gt} = log\, w_{gt} - (x_{gt} - \bar{x}_g)'\beta \tag{3}$$

Thus, the wage index is a *residual* from a log wage equation on gender, race, marital status, education, age, and age-squared. From the perspective of compositional shifts over time explaining earnings developments, the important thing to note here is that even if an individual's characteristics – such as age or occupation – are largely unchanging over time, the average characteristics of subgroup g, x_{gt}, will change with the composition of the group. The second stage involves estimating the first difference specification commonly used in the literature, only at the *group* g and, in our case, region j level:

$$\Delta log\, \hat{w}_{gjt} = \alpha_g + \gamma \Delta log\, u_{gjt} + \epsilon_{jt} \tag{4}$$

In this specification, therefore, γ relates changes in earnings *not* explained by changes in characteristics to changes in the regional unemployment rate. Tables 4 and 5 reports the results from estimating the first-difference specification, following the HSvR wage-indexing approach for the first stage. The results in Table 4 refer to a fixed effects first stage specification; whereas the results in 5 control explicitly for education, occupation and sector. The key result is that in both sets of results the semi-elasticity remains significant and in similar range to earlier estimates -1.64 to -1.85. Also in line with the earlier results, the earnings of incumbents and job changes appear less sensitive to changes in unemployment. There has however, been a significant increase in the coefficient on job changers (albeit statistically insignificant), moving it more in-line with new hires. This tentatively suggests that composition effects are quite important for job changers.

Table 4. Response of earnings to unemployment changes

	First differences specification			
	First stage estimated as fixed effects model			
	(1)	(2)	(3)	(4)
	All workers	Incumbents	Job Changers	New hires
Δ unemployment	-0.498***	-0.204	-2.057	-1.644**
	(0.150)	(0.158)	(1.615)	(0.792)
Observations	42,136	36,029	2,512	3,595

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Observations refer to the number of employees in the first stage.

Table 5. Response of earnings to unemployment changes

	First differences specification			
	First stage controls for education, occupation and sector			
	(1)	(2)	(3)	(4)
	All workers	Incumbents	Job Changers	New hires
Δ unemployment	-0.309*	-0.126	-1.650	-1.853**
	(0.168)	(0.153)	(1.255)	(0.785)
Observations	31,463	27,565	1,550	1,718

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Observations refer to the number of employees in the first stage.

3.1.2 Fixed effects regression for different sub-groups

Next we adopt an approach more akin to GHT, where we try to control for compositional bias by re-estimating the fixed effects specification *within* individual education, age, sector, and occupation groups. This effectively amounts to testing for a higher semi-elasticity for new hires within more homogeneous groups. In all cases, we use quite broad measures of each of the characteristics, mainly to avoid small cell size issues, particularly for new hires, who are only a subset of workers.

The estimated coefficients and semi-elasticities from estimating the fixed effects specification within groups are shown in Tables 6 (age and education) and 7 (occupation and industry groups). The full set of results is summarised in Figure 10, which concentrates on the estimated (total) semi-elasticities (the percentage change in weekly earnings of new hires for a percentage point change in the regional unemployment rate). The key point is that the large, negative, and statistically significant semi-elasticity result holds across all groups (all bars go to the right, i.e., to the negative).

There are two additional results that stand-out. First, there are important differences in the flexibility of earnings of new hires from unemployment by age group. Similar to the the results in HSvR, we find that the age group with the most flexible **earnings** for new hires from unemployment is middle- to older- workers, specifically workers aged 45 and above. We interpret this results as being consistent with the notion that the value of the outside option influences the sensitivity of earnings of new hires from unemployment. Older age groups can neither delay their entry (back) in to the labour market, or, for those close to, but not close *enough* to retirement, they cannot afford to wait out the bad times.

Second, we look at the flexibility of earnings of new hires by education. If the value of the outside option does indeed affect the sensitivity of earnings of new hires from unemployment, then we would expect to observe a lower semi-elasticity for more educated groups, who might be expected to have more attractive outside options. This is indeed what we find. The semi-elasticity for those with a tertiary education is -1.41, compared with -2.66 for those with a primary education. The null hypothesis that the semi-elasticity for primary equals that for tertiary is rejected (p-value 0.049). These results, along with the results in Table 6 would imply that flexibility of earnings of new hires is to some extent endogenous and can depend on (the change of) government policies regarding the level and duration of unemployment benefits during the downturn, openness of the labour market to migration, etc. For instance, one of the reasons why we find stronger results than HSvR might be that the Great Recession in Ireland also caused a major strain on government finances and that this was perceived as a risk that the reliance on unemployment benefits and other forms of social support may not be guaranteed, which in turn lowered the value of unemployed workers staying unemployed.

Figure 10. Semi-elasticities for different groups

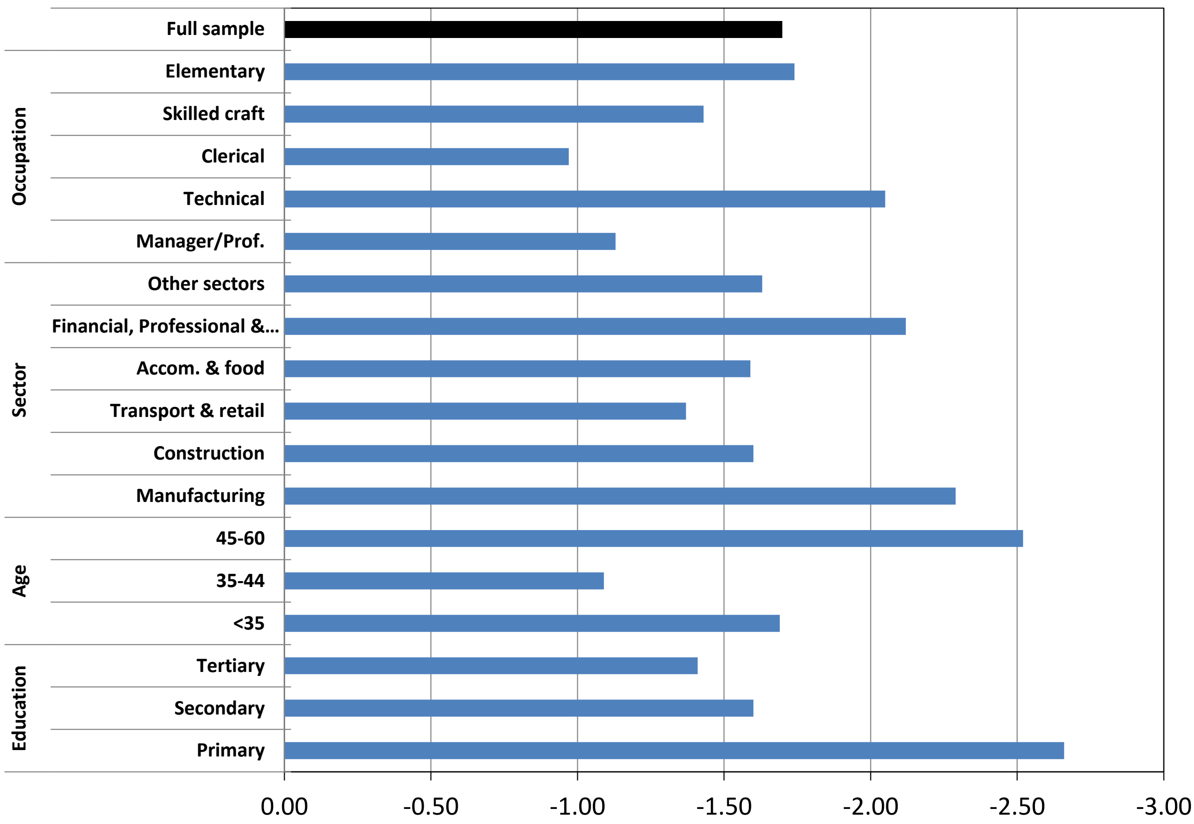

Notes: Based on results shown in Tables 6 and 7.

Table 6. Fixed effects wage regressions by age and education

Log(weekly earnings)	(1)	(2)	(3)
By age group	<35	35-44	45-60
Unemployment (β_u)	-0.855***	-0.277***	-0.006
	(0.114)	(0.116)	(0.155)
Newhire, U	-0.835***	-0.816***	-2.52***
#Unemp ($\beta_{newU,u}$)	(0.211)	(0.328)	(0.384)
Semi-elasticity	-1.69	-1.09	-2.53
Observations	19,523	10,814	10,200
Average # new hires p.a.	268	74	62
By education	Primary	Secondary	Tertiary
Unemployment (β_u)	-0.774	-0.315***	-0.614***
	(0.507)	(0.107)	(0.103)
Newhire, U	-1.894***	-1.285***	-0.799***
#Unemp ($\beta_{newU,u}$)	(0.406)	(0.230)	(0.249)
Semi-elasticity	-2.66	-1.60	-1.41
Observations	6,738	16,640	18,758
Average # new hires p.a.	68	190	154

Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 7. Fixed effects wage regression by occupation (6) and industry (7)

Log(weekly earnings)	(1)	(2)	(3)	(4)	(5)	(6)
By occupation	Manager & prof.	Tech. & assoc. prof.	Clerical & serv.	Skilled agri. craft & plant	Elementary occs.	
Unemployment (β_u)	-0.144	-0.845***	-0.201*	-0.702***	-0.370*	
	(0.113)	(0.238)	(0.120)	(0.180)	(0.224)	
Newhire, U	-1.133***	-1.208*	-0.768***	-0.725*	-1.374***	
#Unemp ($\beta_{newU,u}$)	(0.347)	(0.695)	(0.289)	(0.429)	(0.487)	
Semi-elasticity	-1.13	-2.05	-0.97	-1.43	-1.74	
Observations	12,181	3,051	11,221	4,326	2,472	
Average # new hires p.a.	64	17	86	35	24	
By industry	Manu & industry	Constr.	Transport & retail	Accom. & food	Fin., prof. & admin.	Other sectors
Unemployment (β_u)	-0.373***	-1.602***	-0.361**	-0.537*	-0.894***	-0.904***
	(0.139)	(0.401)	(0.146)	(0.296)	(0.155)	(0.356)
Newhire, U	-1.920***	0.797	-1.011***	-1.052*	-1.230***	-0.723**
#Unemp ($\beta_{newU,u}$)	(0.460)	(1.131)	(0.360)	(0.567)	(0.344)	(0.356)
Semi-elasticity	-2.29	-1.60	-1.37	-1.59	-2.12	-1.63
Average # new hires p.a.	21	7	46	24	43	57

Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

4 Conclusion

This paper analyses the sensitivity of new hires' earnings to changes in the unemployment rate in Ireland over the 2005-14 period. The empirical analysis uses a unique tax database on weekly earnings of over 4,000 workers over a ten-year period. Results from log earnings regressions suggests that a 1% increase in the unemployment rate led to a 1 to 1.7% decline (semi-elasticity) in the weekly pay of new hires, where the lower number comes from the regression which controls for changes in the occupational composition of new hires during the recession. Incumbent workers' earnings are, by comparison, less cyclically sensitive: we find semi-elasticity for incumbents of just 0.37. This finding is robust to different econometric specifications and seems to indicate the presence of significant scarring effects, possibly related to the persistence of bad matches formed during and after the downturn.

Overall, the findings provide robust support that earnings of new hires are most responsive to changes in the economic environment in Ireland and, within this group, the earnings of workers who were previously inactive or unemployed are the most flexible. We find that the weekly earnings of job changers are no more or less responsive to unemployment when compared with workers that do not change jobs. This is the direct opposite of the result in Gertler et al. (2015), who find no negative coefficient on the new hire from unemployment identifier and a large negative coefficient on the new hire from employment coefficient. Even after controlling for compositional shifts, the weekly earnings of new hires are found to be significantly more sensitive to changes in the regional unemployment rate. Our results are in line with the findings in HSvR, but unlike their results, ours are statistically highly significant and robust.

We also find that the sensitivity of earnings of new hires from unemployment is related to age groups and education. In particular, workers who have better education or who have the option to delay entry to the labour market or wait out until retirement show less sensitivity of earnings of new hires from unemployment than other groups. We interpret this as evidence consistent with the notion that the value of the outside option has an effect on wage flexibility of new hires from unemployment. If this is the case, then this flexibility is affected by government policies regarding unemployment benefits as well as the openness of the labour market to migration, which are both likely to have played an important role in Ireland. Our findings suggest that there was a significant weakening in the bargaining power of new hires during the recession, which had a negative impact on their earnings. This may in part be attributable to the abandonment of national wage agreements from 2009/10 onwards which provided some protection for new hires. While the abandonment of these agreements may have prevented an even larger increase in unemployment, our results also suggest that it has increased the wage inequality between workers doing the same job.

References

Abbritti, M. and Fahr, S. (2013). Downward wage rigidity and business cycle asymmetries. *Journal of Monetary Economics*, 60(7):871–886.

Barlevy, G. (2001). Why are the wages of job changers so procyclical? *Journal of Labor Economics*, 19(4):837–878.

Barlevy, G. (2002). The sullying effect of recessions. *The Review of Economic Studies*, 69(1):65–96.

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

Bils, M. J. (1985). Real wages over the business cycle: evidence from panel data. *The Journal of Political Economy*, 93:666–689.

Devereux, P. J. (2001). The cyclicality of real wages within employer-employee matches. *Industrial & Labor Relations Review*, 54(4):835–850.

Fabiani, S., Lamo, A., Messina, J., and Rõõm, T. (2015). European firm adjustment during times of economic crisis. *IZA Journal of Labor Policy*, 4(1):1–28.

Fagan, G. and Messina, J. (2009). Downward wage rigidity and optimal steady-state inflation. *ECB Working Paper No. 1048*.

Gertler, M., Huckfeldt, C., and Trigari, A. (2015). Unemployment fluctuations, match quality and the wage cyclicality of new hires. *Society for Economic Dynamics, 2015 Meeting Papers*.

Glynn, I., Kelly, T., and MacEinri, P. (2013). Irish emigration in an age of austerity. *Cork: Department of Geography, Institute for the Social Sciences in the 21st Century, University College Cork*.

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

Lawless, M., Lydon, R., and McIndoe-Calder, T. (2015). The financial position of irish households. *Quarterly Bulletin Article, Central Bank of Ireland*, pages 66–89.

Leddin, A. (2010). The phillips curve and the wage inflation process in ireland. In Kinsella, S. and Leddin, A., editors, *Understanding Ireland's economic crisis: prospects for recovery*, chapter 7, pages 159–179. Blackhall Press, Dublin.

Linehan, S., Lydon, R., and Scally, J. (2015). Labour cost adjustment during the crisis: Firm-level evidence. *Quarterly Bulletin Articles, Q3 2015*, page 73.

Martins, P. S., Solon, G., and Thomas, J. P. (2012). Measuring what employers do about entry wages over the business cycle: A new approach. *American Economic Journal: Macroeconomics*, 4(4):36–55.

O'Connell, P. J. (1999). *Astonishing success: Economic growth and the labour market in Ireland*. International Labour Organisation, *Employment and Training Papers 44*.

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

Sedláček, P. (2014). Match efficiency and firms' hiring standards. *Journal of Monetary Economics*, 62:123–133.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, pages 25–49.

Solon, G., Barsky, R., and Parker, J. A. (1994). Measuring the cyclicality of real wages: How important is composition bias? *The Quarterly Journal of Economics*, 109(1):1–25.

Topel, R. H. and Ward, M. P. (1992). Job mobility and the careers of young men. *Quarterly Journal of Economics*, 107.

A Additional Charts and Tables (For Online Appendix)

Figure 11. Identification of 'New Hires' in the data

	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014
ID#	10A	10A	10A	10A	10A	10A	10A	10A	10A	10A
Weekly earnings (€)	No	No	No	Yes	Yes	Yes	Yes	No	No	No
New hire from unemployment	0	0	0	1	0	0	0	0	0	0
ID#	11A	11A	11A	11A	11A	11A	11A	11A	11A	11A
Weekly earnings (€)	No	No	No	No	No	No	Yes	Yes	Yes	Yes
New hire from unemployment	0	0	0	0	0	0	1	0	0	0
ID#	12A	12A	12A	12A	12A	12A	12A	12A	12A	12A
Weekly earnings (€)	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
New hire from unemployment	0	1	0	0	0	0	0	0	0	0
ID#	13A	13A	13A	13A	13A	13A	13A	13A	13A	13A
Weekly earnings (€)	Yes	Yes	Yes	No	No	No	No	No	Yes	Yes
New hire from unemployment	0	0	0	0	0	0	0	0	1	0
ID#	14A	14A	14A	14A	14A	14A	14A	14A	14A	14A
Weekly earnings (€)	No	Yes	Yes	Yes	Yes	No	No	No	No	Yes
New hire from unemployment	0	1	0	0	0	0	0	0	0	1
ID#	15A	15A	15A	15A	15A	15A	15A	15A	15A	15A
Weekly earnings (€)	Yes	Yes	No	No	No	No	No	No	No	Yes
New hire from unemployment	0	0	0	0	0	0	0	0	0	1
ID#	16A	16A	16A	16A	16A	16A	16A	16A	16A	16A
Weekly earnings (€)	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes
New hire from unemployment	0	0	0	0	1	0	0	0	0	0
ID#	17A	17A	17A	17A	17A	17A	17A	17A	17A	17A
Weekly earnings (€)	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes
New hire from unemployment	0	0	0	0	0	1	0	0	0	0

Source: HFCS-Admin

FIGURE 12. Identification of 'Job Changers' in the data

	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014
ID#	18A	18A	18A	18A	18A	18A	18A	18A	18A	18A
weekly earnings (€)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Job changer	0	1	0	0	0	0	0	0	0	0
ID#	19A	19A	19A	19A	19A	19A	19A	19A	19A	19A
weekly earnings (€)	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Job changer	0	0	0	0	0	1	0	0	0	0
ID#	20A	20A	20A	20A	20A	20A	20A	20A	20A	20A
Weekly earnings (€)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Job changer	0	0	1	0	0	0	0	0	0	0
ID#	30A	30A	30A	30A	30A	30A	30A	30A	30A	30A
Weekly earnings (€)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Job changer	0	0	0	1	0	0	0	0	0	0
ID#	40A	40A	40A	40A	40A	40A	40A	40A	40A	40A
Weekly earnings (€)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Job changer	0	0	0	0	0	0	0	1	0	0
ID#	50A	50A	50A	50A	50A	50A	50A	50A	50A	50A
Weekly earnings (€)	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes
Job changer	0	0	0	0	0	0	1	0	0	0
ID#	60A	60A	60A	60A	60A	60A	60A	60A	60A	60A
Weekly earnings (€)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Job changer	0	0	0	0	0	0	0	0	1	0
ID#	70A	70A	70A	70A	70A	70A	70A	70A	70A	70A
Weekly earnings (€)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Job changer	0	0	0	0	0	0	0	0	1	0

Source: HFCS-Admin

Figure 13. Average change in real weekly pay following a job change

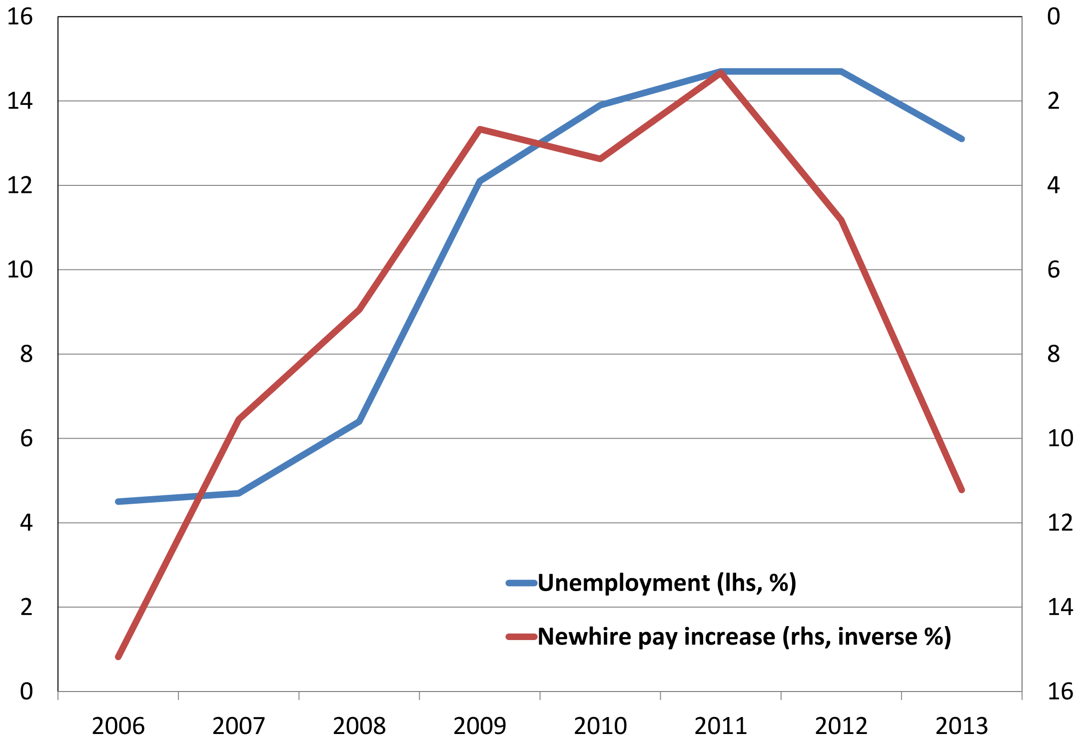

Source: HFCS-Admin

TABLE 8. Difference-in-difference regression log(weekly earnings)

	Incumbents	Job changes	New hires
2006	[Omitted]	-0.0245	-0.0074
		0.0263	0.0177
2007	0.0147	-0.0802	-0.0346
	0.0068	0.0255	0.0173
2008	-0.0010	-0.0127	-0.0432
	0.0065	0.0247	0.0190
2009	0.0000	-0.0297	-0.0849
	0.0063	0.0275	0.0237
2010	-0.0260	-0.0041	-0.0900
	0.0063	0.0234	0.0206
2011	-0.0248	-0.0595	-0.1082
	0.0064	0.0221	0.0198
2012	-0.0364	-0.0189	-0.1380
	0.0066	0.0191	0.0188
2013	-0.0248	0.0151	-0.1260
	0.0066	0.0319	0.0207
2014	[Omitted]		-0.1799
			0.0254

Standard errors in parentheses
Coefficients from a fixed effects regression of log weekly pay on year dummies interacted with new hire and job change dummy variables, including controls for age and age-squared.