

Wage Flexibility and the Great Recession: The Response of the Irish Labour Market

Aedín Doris^{*}, Donal O’Neill^{**} & Olive Sweetman^{*}

NUI Maynooth WP N244-13

November 2013

Abstract

There is considerable debate about the role of wage rigidity in explaining unemployment. Despite a large body of empirical work, no consensus has emerged on the extent of wage rigidity. Previous attempts to empirically examine wage rigidity have been hampered by small samples and measurement error. In this paper we examine nominal wage flexibility in Ireland both in the build up to, and during the Great Recession. The Irish case is particularly interesting because it has been one of the countries most affected by the crisis. Our main analysis is based on earnings data for the entire population of workers in Ireland taken from tax returns, which are free of reporting error. We find a substantial degree of downward wage flexibility in the pre-crisis period. We also observe a significant change in wage dynamics since the crisis began; the proportion of workers receiving wage cuts more than doubled and the proportion receiving wage freezes increased substantially. However, there is considerable heterogeneity in wage changes, with a significant proportion of workers continuing to receive pay rises at the same time as other were receiving pay cuts.

JEL Classification: J31, J38, D31

Keywords: Wage Flexibility, Great Recession

^{**}Corresponding author; National University of Ireland Maynooth and IZA, Bonn. E-mail: donal.oneill@nuim.ie; tel.: 353-1-7083555; fax: 353-1-7083934; address: Rhetoric House, NUI Maynooth, Maynooth, Co. Kildare, Ireland.

^{*}National University of Ireland Maynooth.

We would like to thank Paul Devereux, Uwe Sunde, Frank Walsh and participants at the NERI Labour Market Conference, Dublin, May 2013 for helpful discussions. We would like to thank Ben Berstock, Berni Dunne, John Dunne and Mervyn Ó Luing for help in accessing the Job Churn data and Marion McCann for help with the EU-SILC data.

1. Introduction

The issue of whether wages are rigid or flexible is one that has been central to macroeconomics for many years. Wage stickiness in either direction is of concern as it can explain why nominal shocks translate into real effects.¹ However there is no consensus on the impact of wage rigidity on unemployment. Downward inflexible wages could prevent labour markets from clearing causing unemployment to persist. However there is also a literature that argues that under depression type conditions, with interest rates close to zero, wage flexibility will have little impact on unemployment and may even exacerbate the problem. This paper examines the flexibility of wages in the Irish labour market before and during the Great Recession.

The Irish economy provides a very useful setting for examining the flexibility of wages. Firstly, the Irish labour market is generally held to be flexible with relatively low levels of job protection legislation, above average working time flexibility and above average functional flexibility, measured by the ease with which employers can change the content of work (Andranik 2008). It is therefore interesting to see if flexibility along these dimensions translates into flexibility in wage setting. In the only analysis of Irish microeconomic wage data published to date, nominal rigidity was found to be low by international standards (Dickens et al. (2007)). However the authors themselves expressed misgivings about the quality of the data used and suggested that the Irish results might be due to measurement error. In this paper, we use two data sources for which measurement error is likely to be much less of an issue; the data we use are also more recent.

¹ Such general wage stickiness is usually explained by menu costs such as the management time costs of performance reviews and the administrative costs of adjusting payslip details.

A second reason for interest in wage flexibility in Ireland is that the Irish economy has sustained a serious downturn in recent years. After a period of very rapid growth from 1994 to 2007, when the average annual GDP growth rate was over 7%, the economy collapsed and the average growth rate over 2008-2011 was -1.75%. There are similar patterns for unemployment; having been relatively stable at 4%-5% for most of the early 2000s, the unemployment rate rose from 4.5% in 2007 to 12% in 2009 and continued to rise further to 14.6% in 2011. Inflation averaged 2.53% in the period from 1994 to 2011 but were negative in 2009 (-4.5%) and 2010 (-1%). Given these substantial changes in the macroeconomic environment, it is very interesting to examine the extent to which wages responded during this period.

The predominant explanation for why wages might be downwardly rigid is that employers avoid reducing wages because of the effect on morale. Bewley (1999) examined the downward rigidity of wages in the US during the recession of 1991-1992, and found that managers used wage cuts only in circumstances where the firm faced serious problems. Since the economic crisis in Ireland caused serious problems for many firms, it is plausible that downward nominal wage rigidity would be lessened in these years. In addition, Gordon (1996), in his comment on Akerlof et al.'s paper on the impact of wage rigidity in a low-inflation environment, suggests that nominal wage reductions would no longer be seen as unfair. The fact that inflation dropped and then turned negative during the crisis might also lead us to expect lesser downward nominal wage rigidity.

In this paper we look at nominal wage changes over the pre-crisis and post-crisis periods using a newly-available administrative panel dataset covering the population of Irish workers, known as the 'Job Churn' data. The very large number of observations (700,00-800,000 in the subset of the data that we study) allows for the examination of the wage

change distribution at a level of detail not previously possible. In addition, because these data are derived from employee records returned to the tax authorities by employers, they are free from the reporting errors that plague analyses of survey-based data. We also use the Irish component of the EU-SILC to carry out some supplementary analyses; although based on survey data, the EU-SILC includes additional controls that allow us to examine some possible explanations for the patterns we observe in the Job Churn data.

We find a significant degree of downward wage flexibility in the pre-crisis period in both annual earnings and hourly wages. We also observe a significant response in wage change behaviour since the crisis began; the proportion of workers receiving wage cuts more than doubled and the proportion receiving wage freezes increased substantially. However, there is considerable heterogeneity in wage changes, with a significant proportion of workers continuing to receive pay rises at the same time as other were receiving pay cuts.

2. Literature Review

There is a substantial body of research that uses microdata to examine the extent of wage stickiness. However, as of yet no general consensus has emerged. Much of this work has focused on the US and UK. McLaughlin (1994) analysed PSID data and concluded that wages in the US were flexible; 43% of household heads who did not change employers faced real wage cuts annually, while approximately 17% of the sample faced nominal wage cuts. However, these results have been challenged by a number of authors who argue that the extent of wage cuts in these data may be exaggerated by measurement error. Altonji and Devereux (2000) using both firm level personnel files and household survey data conclude that nominal wage cuts are rare once one accounts for measurement error. More recently Barattieri et al. (2010), using an alternative identification strategy, reach a similar conclusion.

Looking at quarterly SIPP data they find that in a typical quarter in 1996, 48.1% of the survey report a different wage than in previous quarter. However, when adjusted for measurement error this falls to 17.8%.

For the UK, Smith (2000) uses the 1991-1996 British Household Panel Survey (BHPS) to examine wage rigidity. Her initial results indicate that 9% of job stayers experienced zero nominal wage changes from year to year, and that 23% experienced nominal wage reductions. To examine the consequences of measurement error, she uses the fact that the BHPS records whether respondents consulted their pay slips when answering the wage question. On the assumption that measurement error should be lower for those who consult their pay slips than for those who do not, a comparison of these two groups can be used to identify the impact of measurement on wage changes. In contrast to the results for Altonji and Devereux (2000) and Barattieri et al. (2010), she finds that measurement error in household surveys leads to an understatement of the extent of wage flexibility. The proportion of workers reporting no wage change falls from 9% to 5.6% when the sample is restricted to those who consult their payslip. She attributes this difference to rounding error and notes that in contrast to classical measurement error, rounding errors may lead researchers to understate the degree of wage flexibility; for example, a worker whose wage was €10.75 last year and €11.25 this year may round to €11 in both years.

Evidence of wage changes for other countries is more limited. Dickens et al. (2007) report the results of the International Wage Flexibility Project, which analyses individual earnings in 31 different data sets from 16 countries. They find that on average, 8% of workers receive nominal wage freezes, and in many countries wage cuts are rare so that wage change distributions are typically asymmetric. Ireland is unusual in that there is a lower incidence of wage freezes, and almost as many wage cuts are reported as would be if the wage cut

distribution were symmetric. They argue that the data used for the Irish analysis, the European Community Household Panel (ECHP) may explain the unusual Irish results, as it contains fewer observations and more reporting errors than the datasets available for other countries.

Barwell and Schweitzer (2007), Bauer et al. (2007) and Devicienti et al. (2007) use a common methodology that corrects for measurement error in order to identify real and nominal wage rigidities in the U.K., Germany and Italy respectively.² They find far less downward nominal wage rigidity than earlier studies that corrected for measurement error.

Researchers have begun to examine wage adjustment in the Great Recession. Blundell et al. (2013) examine payroll data from the National Employment Survey (NES) data for the UK and find that wages have responded much more to the current recession than to previous recessions. The number of workers experiencing wage freezes has increased from approximately 5% in 1990 to 12% in 2011. However, in line with Smith's (2000) results based on survey data, they find a significant degree of downward wage flexibility in their payroll data; they find that throughout the 1990s and 2000s, almost 20% of stayers report a nominal wage cut.

Daly et al. (2012) argue that downward nominal wage rigidity has been a key reason for the limited extent of real wage reductions in the U.S. in recent years. Elsby et al. (2013) examine wage adjustments in the US in more detail and caution against relying on nominal wage stickiness to explain the high unemployment rates observed during the Great Recession. They report several key features of the wage adjustment process. First, there is always a significant spike at zero in the wage change distribution – between 6% and 20% of workers report exactly the same nominal wage in both years. Secondly, there is always a non-trivial

² A non-technical discussion of the estimator used in these papers is given in Goette et al. (2007).

fraction of workers (between 10% and 20%) who report nominal wage reductions. Thirdly, while the zero-spike increased during the Great Recession, the increase was not substantial, and layoffs were not significantly more prevalent than in earlier severe recessions. Based on their analysis of US data, they suggest that the high unemployment of the Great Recession would have been nearly as high in a world with completely flexible wages. They also analyse UK NES data and find a smaller spike at zero than in the US data, which they attribute to the greater accuracy of the UK payroll-based data. Like Blundell et al. (2013), they find that nominal wage cuts are frequent in the UK.

To our knowledge, apart from the Dickens et al. (2007) paper cited above that used ECHP data from 1994-2001, there has been no study of wage changes in Ireland using microdata. Several recent papers have used a 2007/2008 survey of European firms undertaken for the Wage Dynamics Network to investigate the extent to which wages show downward rigidity and the reasons for this. Du Caju et al. (2013) find that only 2% of firms report having cut wages over the previous five years; the figure for Ireland was just 1%. Babecký et al. (2010) report that 9.6% of firms froze base wages, with a corresponding figure for Ireland of 9%.

Walsh (2012) uses the Earnings, Hours and Employment Costs Survey (EHECS) to examine wage changes during the recession. The EHECS is an employer-based survey that collects information on employment and the firm's total wage bill, and so allows the calculation of average wages in a firm. In addition, it has surveyed firms about the nature of their responses to the recession. Walsh finds that 23% of establishments report cuts in average hourly earnings between 2008 and 2009, rising to 31% between 2009 and 2010.

Bergin et al. (2012) use repeated cross-sectional data from the EHECS and Irish NES³ and find that in the private sector, average earnings and average labour costs increased marginally between 2006 and 2009, while there was no change between 2009 and 2011.

3. Irish Policy Response to the Crisis.

As noted earlier, Ireland was one of the countries worst affected by the Great Recession, with output falling by over 10% in real terms between 2008 and 2010. The effects of the global recession felt elsewhere were compounded in Ireland by the bursting of a property bubble that had inflated during the early 2000s and the subsequent collapse of output and employment in the construction industries. Because bank lending was so highly concentrated in construction activities, Irish banks experienced huge losses following this collapse. On foot of this banking crisis, the Irish government took the decision to guarantee all liabilities of Irish banks in September 2008. However, continued falling tax revenue and exposure to growing bank liabilities resulted in the Irish government deficit going from almost zero in 2008 to 7.4% in 2009, 13.9% in 2010 and a remarkable 30.8% in 2011, when banking losses crystallized. As a result of the outlook for government finances, yields in Irish bonds reached unsustainable levels in 2010, and the government sought and accepted a rescue package from the EU, ECB and IMF.

The crisis resulted in the government undertaking a severe programme of austerity measures, combining tax increases and expenditure cuts. In addition, the government abandoned the national wage setting process that had been in place since 1987, in which unions and participating employers bargained at a national level over wage increases, with tax cuts being offered by the government to encourage wage moderation.⁴ The immediate aim

³ Despite having the same name as the UK dataset, the Irish version does not allow individual workers to be followed over time.

⁴ There is evidence that the nationally-negotiated wage increases came to be regarded as a floor in the private sector during the boom years. As a result, there was a perception that public-sector pay had fallen behind that in

of these measures was to reduce the government deficit. A longer-term aim was to effect an internal devaluation; as a member of the euro area, a nominal exchange rate devaluation was not possible for Ireland, and so only a substantial increase in competitiveness through cuts in labour and other costs could reduce real exchange rates and return the economy to its long-run equilibrium level of output.

Tax increases began in 2009, when the government introduced a new income levy of 1% on incomes up to €100,100 and 2% on income above that. These rates were doubled soon afterwards⁵. 2011 also saw the abolition of the income ceiling on social contributions and a reduction in the upper threshold of the standard tax rate, so that the highest tax rate now applied to income above €32,800 as opposed to €36,400. The standard VAT rate was increased from 21% to 23% in 2012. In addition, a range of new taxes such as a household property tax, a tax on second homes, charges for water usage and a new carbon tax were introduced in an attempt to raise revenue.

At the same time as taxes were increasing, 2010 and 2011 saw cuts in the rate of support provided to most social welfare recipients, especially the young unemployed. In addition the rate of universal child benefit was reduced over successive years, particularly for larger families.

The government also set out to cut payroll costs in the public sector substantially by reducing staff and directly cutting pay. Employment in the public sector was reduced through a major programme of early retirement along with a hiring ban. The number employed fell from 417,600 in the second quarter of 2009 to 377,300 in the second quarter of 2013, a fall of

the private sector, which the government dealt with by commissioning a 'benchmarking' report that was intended to redress this gap. In the benchmarking exercise, public sector employees were given an average pay rise of 8.9%, paid between 2003 and 2005.

⁵ In 2011 both the income levy and the existing health levy were combined into a new tax called the 'Universal Social Charge'.

almost 10%. Pay rates in the public sector were initially reduced via a Pension Levy introduced in 2009 ranging from 5% on incomes of €15,000-€20,000 to 10.5% on earnings above €60,000. Further pay cuts were implemented in 2010,⁶ ranging from 5% on the first €43,000 to 10% on income above €70,000. Most recently, the Haddington Road Agreement (2013) introduced additional wage cuts on higher paid workers, ranging from 5.5% on those earning from €65,000 to €80,000 to 10% on earnings over €185,000. In addition, there were increases in hours worked by all public sector workers, the reduction or elimination of overtime rates, and lower pay scales for new entrants into professions such as teaching. Throughout this period, there has also been severe curtailment of promotions. However, it should be noted that incremental pay increases, laid down in public sector contracts of employment, continued to be paid until 2013, when they were delayed.

Surprisingly given these government implemented pay cuts and the abandonment of national wage bargaining, aggregate data indicate only modest falls in hourly pay rates among public sector workers and stability in the wages of private sector workers since the onset of the crisis (Barrett and McGuinness, 2012). This would suggest that Ireland has been unable to achieve the necessary internal devaluation through wage reductions, perhaps indicating a substantial degree of wage rigidity. However, as acknowledged by the authors, these aggregate data suffer from a number of drawbacks. First it is difficult to control for compositional changes in the workforce that have taken place during the crisis. If workers who have lost their jobs differ from those who continue to be employed then basing average wages only on the population of workers will be misleading. Secondly even if the aggregate wage change is relatively small this may be hiding substantial differences in wage

⁶ These cuts were implemented as part of the Financial Emergency Measures in the Public Interest (No 2) Act 2009, which came into effect on 1 January, 2010.

adjustments across individuals. By following the earnings of individuals over time we address both these issues.

4. Data

Two datasets are used in the analysis. Our main analysis is based on data taken from the Job Churn (JC) dataset, which is an administrative dataset covering the years 2005-2011 that has been compiled by the Central Statistics Office (CSO). The data combines three elements: first data on annual income and weeks worked are provided by the tax authorities (the Revenue Commissioners); the data come from P35 returns, which must be submitted each year by each employer in respect of every worker who was an employee during that year. To these data are added information on workers' age, sex and social welfare class from the Department of Social Welfare. Finally, data from the CSO's own Central Business Register is added to provide information on the sector in which the firms operate and the enterprise's legal form.

There are several strong advantages to using the JC data to examine changes in wages over time. Firstly, because they are administrative data, based on tax returns, they are largely free from measurement error; it is a criminal offence to misreport workers' earnings in these returns. Secondly, the data comprise the entire population of employees in Ireland and so the number of observations is large enough to allow very detailed analysis of job changes; there are up to three million employment records in any year. Thirdly, since employers are obliged to file these returns for every worker, problems associated with non-response and attrition are absent from the data. Finally, the data covers both the period before the crisis (2005-2008) and the period since (2009-2011).

The earnings variable available in the JC data is annual 'reckonable' income for the calendar year; this is gross income after pension contributions have been deducted, as pension contributions are not taxable (up to a limit on the contribution that increases with age). The

disadvantage of such a measure of annual income is that changes in pension contributions by individuals will lead to overstatements of the degree of wage changes, both positive and negative. However, there is a significant advantage to this measure too: it allows us to take into account the Public Sector Pension Levy, mentioned in Section 3 above. Since this levy reduced earnings and entailed no compensating increase in pension entitlements it had the same effect as a reduction in gross pay, but it does not register as such in household surveys that record gross earnings; hence, commentaries on the extent of wage flexibility in Ireland, particularly in the public sector, routinely include a disclaimer that the analysis cannot take into account the Public Sector Pension Levy and therefore understates the true extent of pay cuts. We will be able to take the Pension Levy into account. Unfortunately, the JC data contain no information on hours worked, so it is not possible to distinguish between cuts in hourly pay and cuts in hours worked in these data. For this reason, we supplement the analysis of the JC data with an analysis of another data source, the Irish component of the EU-Survey of Income and Living Conditions (EU-SILC).

EU-SILC data is collected to satisfy a requirement that all members of the European Union collect cross-sectional and longitudinal information on income and living conditions; in Ireland, this requirement is implemented using a dedicated survey. About 5,000 households are interviewed annually. To satisfy the longitudinal requirements, most of the households are re-interviewed for four successive years, with one quarter of the panel being dropped on a rotating basis in any given year, and replaced with new interview households. This means that in the data for any given year, up to three-quarters can be traced back to the previous year. The period covered by the data is the years 2004-2011.

In the Irish data, the ‘income reference year’ for the annual income variable is the 12 months prior to interview, and interviews are carried out on a rolling basis throughout the

year. This means that, for example, the annual income of someone interviewed in July 2012 will in fact refer to seven months of 2012 and five months of 2011; exact matching to a given calendar year is not possible. We have adopted the convention of taking annual income recorded in a given year as referring to that calendar year.

A separate income variable is also recorded in the data that refers to current income; to be precise, it records the income received in the last pay cheque. For this variable, there is no ambiguity about the year referred to. Other important details are also recorded, such as the period covered by the pay cheque, whether it was the usual pay, whether it included overtime payments and weeks, days and hours worked. Respondents were encouraged to consult their pay slips, and whether they did or not was also recorded. Based on this current income variable, an hourly wage variable is also included in the dataset. This is constructed using the information provided on hours worked.

All income variables were subject to careful cleaning by the CSO. This initially involved checking for consistency with the occupation variables provided, but also checking the information on pay and weeks worked and any job changes against the P35 information provided to the Revenue Commissioners by all employers. In a few cases where other documentation was missing, information from P60 forms, which are provided to all employees by employers, was used. Because respondents were encouraged to check their payslips before responding and because the data were subsequently extensively cleaned prior to public release, we expect that reporting error is likely to be less important in the EU-SILC than in the ECHP used by Dickens et al. (2007), even though both are survey data sources.

For both datasets, we focus on job stayers, those who remain with the same employer in successive years. In the JC data, this was based on matching worker identifiers and firm identifiers across pairs of years. Any worker working for the same firm across two calendar

years is therefore identified as a stayer. It should be noted, however, that workers may have changed roles within the firm, and so they may have changed jobs even if they did not change employer. In the EU-SILC data, respondents were asked if they had changed job, and if so, when. Job changes were defined to include promotions at work, so in this case it is possible to identify individuals whose responsibilities had not changed.

For both datasets, we also restrict our samples to workers who had worked for the full year in each pair of years. In the JC data, the weeks worked for the employer is recorded, so we exclude workers who worked for less than 52 weeks in either year. We also exclude all workers who had multiple jobs and all self-employed workers. In the EU-SILC data, respondents were asked about their employment status – including whether any employment was on a part-time or full-time basis – in each of the 12 months prior to the interview, and from this, it was possible to identify individuals who had worked in each month in the income reference year. As noted earlier, there are no hours data available in the JC data and so for ease of comparison, we include both part-time and full-time workers in the EU-SILC sample. The possibility that workers changed hours of work during the observation period is an issue that we return to later.

Because of the different wage-setting mechanisms that pertain in the public and private sectors, we supplement our overall analysis with separate examinations of these sectors. There is no public sector identifier in the JC data. However, the enterprise's NACE code and its legal form can be combined to give a good indication of which sector an individual works in. When defining public sector workers, we omit workers in commercial state enterprises ('semi-state bodies') to the extent possible.

After imposing these restrictions, the number of observations remaining lies between 700,000 and 800,000 in the JC data, and between 800 and 1,700 in the EU-SILC data.

5. Results

a. Analysis of Job Churn Data

To analyze wage dynamics, we first look at annual earnings changes in the JC data for each pair of years between 2005 and 2011. Following Ziliak et al. (2011), we calculate percentage earnings changes using the arc percent change method. In particular the percentage change in earnings is measured as $\frac{y_{it} - y_{it-1}}{y_{it-1}}$, where y_{it} is earnings for person i in time t and y_{it-1} is earnings for person i in time $t-1$. The key advantage of this method is that it is symmetric in gains and losses.

Table 1 presents descriptive statistics of the annual earnings changes from 2005/2006 to 2010/2011. Column 2 shows the median earnings changes for each pair of years. The growth in earnings in the pre-crisis period is evident in the numbers reported for 2005/2006 and 2007/2008, with median growth rates of between 4.5% and 6.1%. The impact of the crisis is clearly observed in the later period, with median wage *reductions* of about 1% in 2008/2009 and 2009/2010. These wage changes are consistent with the relatively small changes in average earnings reported by Barrett and McGuinness (2012). However, as noted earlier, aggregate measures may hide important differences across the distribution. Columns 3-5 report the proportion of workers receiving an earnings freeze, an earnings cut and an earnings rise. Similar to Blundell et al. (2013), we classify a change of less than 0.1% as an earnings freeze. These data reveal substantial flexibility. In the pre-crisis period, we find that between 17% and 23% of workers experience earnings cuts. These findings are similar to the UK findings of Blundell et al. (2013) who found, using payroll data, that during the 1990s and 2000s almost 20% of job stayers report a nominal wage cut.

Not surprisingly, during this period in Ireland between 74% and 80% of employees received earnings increases. However, this pattern changed dramatically with the onset of the

crisis. The proportion of workers experiencing earnings cuts increased to more than 50% in both 2008/2009 and 2009/2010. This figure fell slightly in 2010/2011, but was still high at 39%. While these figures illustrate a significant wage response to the crisis, it is important to note that while many workers were having their earnings cut in recent years, a substantial proportion of workers experienced earnings increases. In fact, the proportion of workers who experienced increased earnings never fell below 40% over this period.

Dickens et al. (2007) propose a simple measure of downward nominal wage rigidity, which is based on the assumption that everyone who had a nominal wage freeze would have had a nominal wage cut in the absence of any rigidity. The measure is defined as the ratio of the proportion of workers receiving cuts to the proportion receiving either cuts or freezes. In Dickens et al.'s cross country comparison, they found that the average degree of downward rigidity across years and datasets was 28%. Their average for Ireland was 4%, the lowest of all countries covered, but as noted earlier, they expressed reservations about the Irish data. We have calculated the same measure and report our findings in the final column of Table 1. These figures show that rigidity was about 11%-12% in the pre-crisis period; although these figures are higher than that reported by Dickens et al. for Ireland, they are comparable to the numbers reported for Denmark and France, which are the two most flexible countries after Ireland. At the onset of the crisis, measured downward rigidity fell substantially to 6% in 2008/2009 and 7.4% in 2009/2010, and then rose again to 14.7% in 2010/2011.

To look at these earnings changes in more detail, Figure 1 shows the histograms of annual earnings changes in each of the years. The very large sample sizes in the JC data allow us to describe the distribution of earnings dynamics in great detail. The red line in each histogram indicates an earnings freeze, defined as a change in annual earnings of less than 0.1%. These histograms display many of the features of wage dynamics noted by Elsby et al.

(2013) in their discussion of US wages. Firstly, in each year we note a significant spike in the nominal earnings change distribution at zero. Secondly, there is always a non-trivial fraction of workers who report nominal earnings reductions. This could reflect changes to overtime rates or reductions in hours worked as well as changes in hourly rates; we will return to this issue later in the paper. Thirdly, the spike at zero increased during the Great Recession. However, in contrast to Elsby et al. (2013), the increase in the spike at zero for Irish workers during the Great Recession is much more dramatic than in the US, with the height of the spike doubling in Ireland between 2007/2008 and 2010/2011. Furthermore, the main increase in the spike occurred in 2010/2011, following two years of substantial earnings cuts. We might have expected that if wage rigidity were a strong feature of the wage setting process, the strongest increase in the zero spike would have occurred in the early years of the crisis. In future research, it will be interesting to examine whether those workers receiving an earnings freeze in later years were the same workers who had previously experienced significant cuts in earnings.

Another notable feature of these histograms is that the increased spike at zero was also accompanied by a substantial increase in the proportion of workers receiving pay cuts. Interestingly, for those affected by pay cuts, the median cut in each of the years from 2005/2006 to 2009/2010 was relatively stable at 5%-6%, falling somewhat in 2010/2011 to 3.7%.

A major focus of policy discussion during the crisis centred on the relative wage adjustments in the public and private sectors. Table 2 presents earnings dynamics separately for public and private sector workers. Looking at the pre-crisis years, we see that the earnings behaviour was relatively similar in both sectors; between 17% and 23% of workers in both sectors received pay cuts, with 74%-84% receiving pay increases. A comparison of the

experiences of workers in both sectors during the crisis is of particular interest. A key government response to the crisis involved reform of the public sector, including the imposition of a series of direct pay cuts on public sector workers. These cuts are evident in 2008/2009 and 2009/2010. As noted earlier, the income measure in these data allows us to see the effect of the Public Sector Pension Levy in 2009; 62% of public sector workers received a pay-cut in the 2008/2009 period. In 2009/2010 the number of public sector workers experiencing a pay cut increased to 82%, reflecting direct pay cuts. The median worker in the public sector experienced a 6% reduction in earnings. Although there were no legislated pay cuts in 2010/2011, 36% of public sector workers received a reduction in annual earnings.

The earnings change distributions for private sector workers also reveal a significant response to the crisis. The proportion experiencing pay cuts increased from 23% in the year immediately preceding the crisis to almost 50% in 2008/2009 and 2009/2010. The figure for 2010/2011 remained high at 40%. In 2009/2010, the median earnings change in the private sector was zero. However, this masks the fact that 47% of private sector workers received a pay cut in this period, while about the same proportion received pay increases. This analysis clearly reveals the dynamic and heterogeneous response of earnings in the Irish labour market during the crisis, a response that can be seen only by looking at this type of data.

The histograms reported in Figures 2 and 3 for public and private sector workers respectively show these earnings dynamics in more detail. Looking at Figure 2, we see the clear shift to the left of the earnings change distribution for public sector workers, representing the substantial cuts in earnings for workers in this sector over the crisis period. We also see the emergence of a strong spike at zero for public sector workers in 2010/2011

which was entirely absent for these workers prior to the crisis.⁷ The histograms for private sector workers are in keeping with the discussion for all workers; a persistent spike at zero that increased dramatically during the crisis, combined with a substantial increase in the proportion of workers receiving pay cuts at the height of the crisis.

b. Analysis of EU-SILC Data

As mentioned earlier, the JC data comprises the population of workers and the information is taken from employers' tax forms, so has the advantage that it is unlikely to be subject to measurement error. However a disadvantage of these data is that the income period is annual and therefore changes in annual earnings may reflect changes in hours worked as well as changes in the rates of pay. To examine this in more detail, we look at the EU-SILC data. These data are survey-based and have much smaller sample sizes, but they do contain information on hours worked, allowing the construction of an hourly pay rate. A comparison of wage changes using EU-SILC and JC will be useful in helping understand the dynamics presented earlier.

We begin by comparing the dynamics for annual earnings presented for the JC data reported above with similar summary measures based on the EU-SILC data. As discussed above, the restrictions on the EU-SILC sample are similar to those used for the JC analysis. We examine gross annual earnings from the EU-SILC data.

Table 3 shows the results from EU-SILC, alongside the results from JC, reproduced from Table 1. Looking first at the median changes, we see that with the exception of 2008/2009, both data sets give very similar aggregate results. In both data sets, earnings growth fell from about 6% in 2005/2006 to about 4.5% in 2007/2008. Furthermore, both data

⁷ There are, however, substantial spikes away from zero; there was strict adherence to the national wage agreements in the public sector, so these spikes are likely to correspond to the wage rises from these agreements.

sets indicate negative earnings growth in 2009/2010 and very small positive earnings growth in 2010/2011. The only difference between the two series arises in 2008/2009, where the EU-SILC earnings data shows that the median rise in earnings was 5% in 2008/2009, while the JC data reports a median fall of 0.6%. This can be explained by the fact that, as mentioned above, 2009 was the year in which the Public Sector Pension Levy took effect. This levy did not reduce the measure of earnings used in the EU-SILC (gross earnings), but did reduce the measure of earnings used in the JC data (taxable earnings).

The trends in the proportions receiving earnings freezes, cuts or increases are also similar across the two data sets. Both data sets show a rise in the proportion receiving earnings cuts and a fall in the proportion receiving earnings increases during the crisis. In both datasets, about 6% of workers receive a pay freeze in 2010/2011 compared to 2.5% in 2005/2006. As well as the difference in the 2008/2009 figures previously discussed, other year pairs do show differences in the *levels* of pay cuts and pay rises. However, in many years, the differences are small. For example, in both datasets the proportion of workers receiving an earnings cut reached a high of over 50% in 2009/2010 as the effect of the recession hit Irish workers hardest, before falling to just under 40% in 2010/2011. The heterogeneity in earnings responses revealed above in the JC data is also evident in the EU-SILC data; while many workers were receiving earnings cuts during the crisis, a substantial proportion (over 40%) continued to receive earnings increases.

It has been well documented (e.g. Walsh, 2012) that firms in Ireland responded to the crisis in part by adjusting hours of work, which would be reflected in changes in annual earnings with no corresponding change in hourly pay. The primary advantage of using the EU-SILC data is that it provides data on hours worked and therefore allows us to examine dynamics in hourly pay. The results are given in Table 4; to allow comparison, we reproduce

the EU-SILC results on annual earnings from Table 3 in the first five columns. The most notable difference between the two series is the significant increase in workers who reported a pay freeze when using hourly pay as opposed to annual pay; over 12% of workers reported hourly pay freezes in 2010/2011 compared to 6% with annual freezes. Despite this difference, the major features of wage dynamics reported earlier for annual earnings using both the JC and EU-SILC data are still very evident when we use hourly pay. The proportion of workers receiving a cut in hourly pay increased from below 30% in 2004/2005 to almost 50% in 2009/2010, as the labour market reacted significantly to the crisis. As before the cuts experienced by many workers are hidden in aggregate data by the fact that at the same time, at the height of the crisis, a significant proportion of workers continued to receive increases in hourly pay.

Since the EU-SILC data are survey-based, they are more likely to suffer from measurement error than the JC data.⁸ To examine the likely impact of measurement error on the EU-SILC findings, we follow Smith (2000). She examined the impact of measurement error on wage dynamics in the BHPS using the fact that respondents in the BHPS were told that they could consult their pay slips when answering the wage question. On the assumption that measurement error should be lower for those who consult their pay slips, a comparison of the two groups illustrates the impact of measurement error on wage changes. The EU-SILC data allows us to conduct a similar comparison. The results are given in Table 5. Our findings are similar to Smith (2000) in that we find that the proportion of workers reporting wage freezes is smaller when the sample is restricted to those workers who consulted their pay slip. As noted earlier, this is not consistent with classical measurement error but is consistent with rounding of reported earnings. While there are some differences in the levels of freezes, cuts and increases across the two samples, the broad features of wage dynamics highlighted

⁸ As noted earlier, the EU-SILC earnings data is subject to extensive cleaning prior to public release.

throughout this analysis remain evident when we restrict our attention to workers who consulted their wage slips. These results suggest that measurement error is not the driving force behind our findings.

6. Conclusion

A large body of macroeconomic research emphasizes the role of wage rigidity in accounting for unemployment. Excess supply in any market is typically eliminated by price reductions. Downward inflexible wages could prevent labour markets from clearing, causing unemployment to persist. However there is also a literature that argues that under depression-type conditions, with interest rates close to zero, wage flexibility will have little impact on unemployment and may exacerbate the problem.

In this paper we look at nominal wage flexibility in Ireland both before and during the Great Recession. Previous attempts to measure wage rigidity have been hampered by small samples and measurement error. Our main analysis is based on earnings data for the entire population of workers in Ireland taken from tax returns filed by their employers. Since it is a criminal offense to misreport taxable earnings, we are confident that these data are free of the misreporting that can plague survey data. We also use a supplementary dataset in order to account for changes in hours worked.

We find a significant degree of downward wage flexibility in the pre-crisis period in both annual earnings and hourly wages. We also observe a significant response in wage change behaviour since the crisis began. The proportion of workers receiving wage cuts more than doubles and the spike at zero increases substantially, particularly in 2010/2011. However, there is substantial heterogeneity of wage changes, with a significant proportion of workers continuing to receive pay rises at the same time as other were receiving pay cuts.

Our analysis confirms previous research that Ireland's labour market is a flexible one. It is important to take this into account when devising policies to address the severe unemployment crisis in Ireland.

References

- AKERLOF, G. A., W. T. DICKENS, and G. L. PERRY (1996): "The Macroeconomics of Low Inflation," *Brookings Papers on Economic Activity*, 1, 1-76.
- ALTONJI, J., and P. DEVEREUX (2000): "The Extent and Consequences of Downward Nominal Wage Rigidity," *Research in Labor Economics*, 19, 383-431.
- ANDRANIK, T. (2008): "On the European Readiness for Flexicurity: Empirical Evidence with OECD/HBS Methodologies and Reform Proposals." WIS Discussion paper no. 160.
- BABECKÝ, J., P. D. CAJU, T. KOSMA, M. LAWLESS, J. MESSINA, and T. RÕÖM (2010): "Downward Nominal and Real Wage Rigidity: Survey Evidence from European Firms," *The Scandinavian Journal of Economics*, 112, 884-910.
- BARATTIERI, A., S. BASU, and P. GOTTSCHALK (2010): "Some Evidence on the Importance of Sticky Wages," Cambridge, Mass: National Bureau of Economic Research.
- BARRETT, A., and S. MCGUINNESS (2012): "The Irish Labour Market and the Great Recession," *CESifo DICE Report*, 10, 27-33.
- BARWELL, R. D., and M. E. SCHWEITZER (2007): "The Incidence of Nominal and Real Wage Rigidities in Great Britain: 1978-98," *The Economic Journal*, 117, F553-F569.
- BAUER, T., H. BONIN, L. GOETTE, and U. SUNDE (2007): "Real and Nominal Wage Rigidities and the Rate of Inflation: Evidence from West German Micro Data," *The Economic Journal*, 117, F508-F529.
- BERGIN, A., E. KELLY, and MCGUINNESS (2012): "Explaining Changes in Earnings and Labour Costs During the Recession," Dublin: ESRI.
- BEWLEY, T. (1999): *Why Wages Don't Fall During a Recession*. Cambridge, MA: Harvard University Press.
- BLUNDELL, R., C. CRAWFORD, and J. WENCHAO (2013): "What Can Wages and Employment Tell Us About the UK's Productivity Puzzle." Working Paper No. 13/11, Institute for Fiscal Studies.
- DALY, M., B. HOBIJN, and B. LUCKING (2012): "Why Has Wage Growth Stayed Strong?," *Federal Reserve Bank of San Francisco Economic Letter*.
- DEVICIENTI, F., A. MAIDA, and P. SESTITO (2007): "Downward Wage Rigidity in Italy: Micro-Based Measures and Implications," *The Economic Journal*, 117, F530-F552.
- DICKENS, W. T., L. GOETTE, E. L. GROSHEN, S. HOLDEN, J. MESSINA, M. E. SCHWEITZER, J. TURUNEN, and M. E. WARD (2007): "How Wages Change: Micro Evidence from the International Wage Flexibility Project," *Journal of Economic Perspectives*, 21, 195-214.
- DRUANT, M., S. FABIANI, G. KEZDI, A. LAMO, F. MARTINS, and R. SABBATINI (2009): "How Are Firms' Wages and Prices Linked: Survey Evidence in Europe."
- DU CAJU, P., T. KOSMA, M. LAWLESS, J. MESSINA, and T. ROOM (2013): "Why Firms Avoid Cutting Wages: Survey Evidence Form European Firms."
- ELSBY, M. W., D. SHIN, and G. SOLON (2013): "Wage Adjustment in the Great Recession," Cambridge, Mass.
- GOETTE, L., U. SUNDE, and T. BAUER (2007): "Wage Rigidity: Measurement, Causes and Consequences," *The Economic Journal*, 117, F499-F507.
- GORDON, R. J. (1996): "Comment on Akerlof, Dickens and Perry(1996). The Macroeconomics of Low Inflation," *Brookings Paper on Economic Activity*, 1, 60-66.
- MCLAUGHLIN, K. J. (1994): "Rigid Wages?," *Journal of Monetary Economics*, 34, 383-414.
- SMITH, J. C. (2000): "Nominal Wage Rigidity in the United Kingdom," *The Economic Journal*, 110, C176-C195.
- WALSH, K. (2011/2012): "Wage Bill Change in Ireland During Recession - How Have Employers Reacted to the Downturn?," *Journal of Statistical and Social Inquiry of*

Ireland, XLI, 39-70.

ZILIAK, J. P., B. HARDY, and C. BOLLINGER (2011): "Earnings Volatility in America: Evidence from Matched Cps," *Labour Economics*, 18, 742-754.

Table 1: Earnings Dynamics, Job Churn Data

All Job Stayers

	Median Change (2)	% Freezes (3)	% Cuts (4)	% Increases (5)	Nominal Wage Rigidity (3)/((3)+(4))
2005/2006	0.060	.025	.172	.804	.126
2006/2007	.061	.025	.176	.799	.124
2007/2008	.045	.028	.229	.742	.110
2008/2009	-.006	.033	.527	.440	.060
2009/2010	-.011	.044	.552	.403	.074
2010/2011	.006	.068	.393	.539	.147

Table 2: Earnings Dynamics, Job Churn Data

Public and Private Job Stayers

	All				Public Sector				Private Sector			
	Median Change	% Freezes	% Cuts	% Increases	Median Change	% Freezes	% Cuts	% Increases	Median Change	% Freezes	% Cuts	% Increases
2005/2006	0.060	.025	.172	.804	.061	.005	.171	.824	.059	.030	.172	.798
2006/2007	.061	.025	.176	.799	.068	.005	.153	.843	.059	.030	.180	.791
2007/2008	.045	.028	.229	.742	.046	.006	.218	.775	.044	.034	.231	.735
2008/2009	-.006	.033	.527	.440	-.019	.013	.617	.375	-.001	.041	.501	.458
2009/2010	-.011	.044	.552	.403	-.061	.010	.816	.182	0	.057	.471	.472
2010/2011	.006	.068	.393	.539	.010	.048	.363	.589	.004	.078	.399	.522

Table 3: Earnings Dynamics, EU-SILC and Job Churn Data**All Job Stayers**

Nominal Annual Earnings EU SILC						Nominal Annual Earnings Job Churn (Taken from Table 1)				
	N	Median Change	% Freezes	% Cuts	% Increases		Median Change	% Freezes	% Cuts	% Increases
2004-2005	1599	.067	.026	.278	.696					
2005-2006	1705	.063	.025	.273	.703		0.060	.025	.172	.804
2006-2007	1567	.069	.028	.234	.738		.061	.025	.176	.799
2007-2008	1490	.047	.036	.283	.680		.045	.028	.229	.742
2008-2009	1294	.051	.009	.342	.650		-.006	.033	.527	.440
2009-2010	1254	-.003	.013	.504	.482		-.011	.044	.552	.403
2010-2011	863	.009	.058	.394	.548		.006	.068	.393	.539

Table 4: Earnings and Hourly Wage Dynamics, EU-SILC**All Job Stayers**

Nominal Annual Earnings						Nominal Hourly Wages				
	N	Median Change	% Freezes	% Cuts	% Increases	N	Median change	% Freezes	% Cuts	% Increases
2004-2005	1599	.067	.026	.278	.696	1558	.061	.051	.283	.676
2005-2006	1705	.063	.025	.273	.703	1668	.061	.046	.277	.678
2006-2007	1567	.069	.028	.234	.738	1540	.060	.037	.275	.688
2007-2008	1490	.047	.036	.283	.680	1484	.049	.051	.290	.659
2008-2009	1294	.051	.009	.342	.650	1267	.025	.069	.361	.570
2009-2010	1254	-.003	.013	.504	.482	1214	0	.086	.479	.436
2010-2011	863	.009	.058	.394	.548	835	0	.122	.456	.423

Table 5: Earnings Dynamics, EU-SILC**Job Stayers with Pay Slips and All Job Stayers**

	Pay Slips Available					All (Taken from Table 3)				
	N	Median Change	% Freezes	% Cuts	% Increases	N	Median Change	% Freezes	% Cuts	% Increases
2006-2007	352	.064	.026	.216	.759	1567	.069	.028	.234	.738
2007-2008	830	.049	.016	.270	.714	1490	.047	.036	.283	.680
2008-2009	690	.052	.006	.343	.651	1294	.051	.009	.342	.650
2009-2010	651	-.015	.005	.539	.456	1254	-.003	.013	.504	.482
2010-2011	434	.011	.039	.387	.573	863	.009	.058	.394	.548

Figure 1: Earnings Dynamics, Job Churn Data: All Job Stayers




Figure 2: Earnings Dynamics, Job Churn Data: Public Sector Job Stayers




Figure 3: Earnings Dynamics, Job Churn Data: Private Sector Job Stayers

