

How Important are Firms in Explaining Wage Changes During a Recession?

By AEDÍN DORIS[†], DONAL O'NEILL[‡] and OLIVE SWEETMAN[†]

[†]*Maynooth University*

[‡]*Maynooth University and IZA*

Final version received 7 January 2019.

During the Great Recession, many Irish workers experienced nominal earnings reductions, with about 50% of private sector employees receiving pay cuts at the height of the crisis. However, at the same time, a substantial minority of workers continued to receive pay increases. In this paper we use a unique dataset containing earnings on every worker in Ireland to examine the relative roles of worker and firm characteristics in explaining this heterogeneity in earnings dynamics. Our results show that between-firm effects play a small role in determining pay changes in Ireland. Although between-firm effects became more important in the peak year of the economic crisis, the vast majority of earnings changes continued to be driven by within-firm forces. These findings raise a number of important questions about the role of morale and fairness in the wage-setting process.

INTRODUCTION

In an international study covering sixteen countries, Dickens *et al.* (2007) describe as 'remarkable' the variation in percentage wage changes across individuals in nearly every country and in every year. This heterogeneity was also evident more recently, during the Great Recession. Elsby *et al.* (2016) find substantial heterogeneity in wage changes over this period in the USA and the UK. Doris *et al.* (2015), analysing wage adjustment in Ireland from 2005 to 2013, find that 50% of private sector job stayers received pay cuts at the height of the recession, while a substantial minority of workers continued to receive earnings increases. In this paper we examine whether these pay cuts were concentrated in particular firms or whether all firms were cutting wages for some workers but not for others.

The role of the firm in explaining the variation of wage changes across workers is important because it can shed some light on reward structures within firms. A compressed distribution of pay changes within firms may reduce effort incentives. On the other hand, substantial variation in within-firm pay changes that are not obviously related to productivity may have a negative impact on perceptions of fairness. Such fairness considerations are likely to be more important in times of recession, when firms may be considering pay cuts. The seminal work by Bewley (1999) on why wages do not fall during a recession finds that firms were very reluctant to cut nominal pay, and that almost all the firms that had cut wages in his study did so due to financial distress; none of the cuts was due to a weak labour market (Bewley 1999, p. 200). The main reason given by firms for their reluctance to cut pay was that the negative effects on morale made cuts counterproductive. This would lead us to expect that some firms—those in real financial difficulty—would be willing to cut pay during the Great Recession, but others would avoid cuts altogether. On this basis, we might expect to see the relative importance of incentives and fairness changing when economic conditions worsen; the between-firm component of wage changes is likely to grow substantially with the onset of a recession, with a corresponding reduction in within-firm variation. On the other hand, if workers

within firms are segregated, perhaps on the basis of occupations across which pay comparisons are unlikely, then it may be possible to cut wages for a subset of workers without giving rise to morale effects. In this case a recession that particularly affects a subset of workers may lead to a rise in within-firm variation in pay cuts.

The increasing availability of matched employer–employee datasets has allowed researchers to examine the role of firms in the wage-setting process (Abowd *et al.* 1999; Cardoso 1997, 1999, 2000; Dustmann and Meghir 2005; Lazear and Shaw 2009a; Sørensen and Vejlín 2011; Card *et al.* 2013). Typically, these studies find a significant role for both worker- and firm-level characteristics in wage levels, but a stronger role for individual characteristics in determining wage changes.

In this paper, we examine wage changes in Ireland from 2005 to 2013, a period that includes the years of the Great Recession. During this time, the Irish labour market was under severe stress, with unemployment increasing from 5.0% in January 2008 to a peak of 15.2% in January 2012. We use an administrative matched employer–employee dataset covering every firm and worker in Ireland. The availability of administrative data reduces the potential biases from measurement error, which is often an important problem in examining earnings changes. In addition, the fact that the data cover the population of both firms and workers removes problems associated with small sample sizes and selectivity in examining pay dispersion within firms. We apply a variance decomposition technique to examine the relative contributions of within-firm and between-firm heterogeneity to the overall observed heterogeneity in earnings changes for job stayers documented in Doris *et al.* (2015). Following this, we examine the role of available observable characteristics in determining individuals' earnings changes within firms, allowing for differences by firm size.

Our results show that most of the variation in pay changes in Ireland occurs within firms, indicating that firms have considerable latitude in cutting wages. Although between-firm effects become more important in the peak year of the economic crisis, even then, at the height of the crisis, the vast majority of earnings changes are driven by individual worker characteristics. As well as being informative about the wage-setting process, these findings also raise a number of questions about the role of morale during recessions.

The remainder of the paper proceeds as follows. Section I discusses the macroeconomic and policy context for the period covered by our data. Section II reviews the relevant literature. Section III describes the data used for our analysis. Section IV outlines the statistical and econometric methodologies used. Section V presents the results, and Section VI concludes.

I. THE IRISH MACROECONOMIC AND POLICY CONTEXT, 2005–2013

As mentioned earlier, Ireland was one of the countries worst affected by the Great Recession. After a period of very rapid growth from 1994 to 2007, when the average annual GDP growth rate was over 7%, the economy collapsed, with output falling by over 10% in real terms between 2008 and 2010. Labour market indicators followed this pattern of boom and bust. Employment grew by 46% between 1998 and 2007, but then fell by 15.9% between 2008 and 2012. Similarly, the unemployment rate, which had been relatively stable at 4–5% for most of the 2000s, rose from 5.0% in January 2008 to 13.2% in January 2010, and continued to rise further to a peak of 15.2% in January 2012. It had fallen back to 12.2% by the end of 2013. It is worth noting that the rise in unemployment was particularly evident for males. The male unemployment rate increased from 4.8% in

mid-2007 to 18.2% in early 2012; over the same period, the female unemployment rate increased from 4.5% to 11.5%.

The effects of the global recession felt elsewhere were compounded in Ireland by the bursting of a property bubble and the subsequent collapse of output and employment in construction-related sectors. Because bank lending was so highly concentrated in construction, Irish banks experienced huge losses and the government decided to guarantee all bank liabilities in 2008. However, continued falling tax revenue and exposure to bank liabilities resulted in the government deficit rising from almost zero in 2007 to 13.8% of GDP in 2009 and a remarkable 32.1% in 2010 due to the banking bailout.

The crisis resulted in the government undertaking a severe programme of austerity measures, combining tax increases and expenditure cuts. As part of the expenditure cuts, three rounds of pay cuts were undertaken in the public sector between 2009 and 2013, with cuts ranging from 5% to 10.5% at each round. Although public sector workers are not the focus of this paper, it might be expected that public sector pay cuts would have a demonstration effect that would make similar cuts more acceptable for private sector workers. An additional factor that may have made nominal wage cuts more acceptable was that inflation was negative in 2009 (-4.5%) and 2010 (-1%). Indeed, Doris *et al.* (2015) find that the proportion of private sector job stayers receiving nominal earnings cuts almost trebled during the crisis, rising from 18% in 2006 to 50% at the height of the crisis in 2009. The prevalence of both pay cuts and job losses was widely discussed in the media, so workers were aware of the dramatic changes in the labour market as they occurred.

It is worth commenting on the institutional framework for wage-setting that prevailed in Ireland, and how it changed during the crisis. From 1987 to 2009, Ireland operated a system of national pay agreements. Under this system, unions, government and participating employers bargained at a national level over wage increases. These agreements specified a schedule of percentage pay increases to be implemented over three-year periods, as well as other non-pay measures. While these agreements were always rigidly adhered to in the public sector, for private sector employers, adherence was effectively voluntary; by the start of our sample period, anecdotal evidence suggests that many firms granted wage rises in excess of the negotiated rates. In 2009, the national wage-setting process was abandoned, having been put under pressure by austerity measures undertaken by the government in order to restore public finances. On the collapse of the agreements, a period of enterprise-level bargaining ensued. Under Irish employment legislation, pay cuts constitute a change to the contract of employment and as such must be agreed by the worker. During the crisis period in Ireland, there seems to have been little difficulty in securing such agreement, given the prevalence of pay cuts noted above and the absence of public protests.

The data that we analyse in this paper cover the last few years of the Irish boom, the sudden onset of the economic crisis, the ongoing period of the Great Recession and the first year of the recovery of the labour market. Having access to population earnings data on all firms over this period provides a rare opportunity to investigate how dramatic labour market changes impact on firms and their workers.

II. LITERATURE REVIEW

Initial analyses of the role of firms in the wage-setting process tended to focus on wages within a single firm. For instance, Baker *et al.* (1994) analyse salary data for a single firm

from 1969 to 1988. They find significant variation in both wage levels and wage growth within the firm, even when the analysis is restricted to workers at the same grade. They propose a simple learning model based on unobserved ability to explain the observed patterns.

More detailed comparative analysis of firms' wage-setting policies have until recently been limited by a lack of appropriate data. Abowd *et al.* (1999) provide an early analysis of the role of firms in wage-setting using matched employer–employee data. Their focus is on earnings levels rather than changes, and their data cover over one million French workers from about 50,000 firms. They find that firm effects, while important, are not as important as individual effects in explaining the variation in wage levels in France. More recently, Lazear and Shaw (2009a) provide a collection of country studies analysing the role of the firm in wage-setting. They find that countries are very similar in the structure of their wage levels, with the within-firm and between-firm variations both contributing strongly to the overall standard deviation, with slightly less of the variation coming from wage differences between firms.

Barth *et al.* (2016) focus on the role of the firm in the rise of earnings inequality in the USA in recent decades. They find that the share of the variance of earnings accounted for by the between-firm component grew from 46% to 49% from 1992 to 2007, and about two-thirds of the growth in the overall variance of earnings between those two years was accounted for by between-firm effects. Results for the USA shown in Song *et al.* (2019) indicate a between-firm component somewhat lower than in Barth *et al.* (2016). Alvarez *et al.* (2018) also investigate a change in earnings inequality—this time, a substantial decrease observed in Brazil between 1996 and 2012. Their results show that the between-firm component of earnings inequality was higher than in the USA between 1988 and 2012, accounting for around 55–60% of total inequality. They find that the fall in between-firm inequality contributed more to the fall in overall inequality over the whole period. Finally, Cardoso (1997) examines both worker and employer attributes as sources of wage dispersion and of its rising trend in Portugal during the 1980s and early 1990s, and concludes that over 60% of the total dispersion in hourly wages is accounted for by between-firm variation.

Many of the above-mentioned papers focus on the extent to which within- and between-firm inequalities contribute to changes in inequality in earnings levels, but less research has been done on the importance of the firm in explaining the variance of earnings *changes*, which is the focus of this paper. However, several of the papers in the Lazear and Shaw (2009a) volume focus on wage growth as well as levels. For the Netherlands, Borghans and Kriechel (2009) find that wage growth differences within firms account for about 87% of the variation in individual wage changes. They also investigate how these results varied by firm size. They find that the overall variance of wage changes was very similar across firm sizes, but that the magnitudes of the components differed substantially. In particular, the firm-level component was much higher (36%) in smaller firms and lower (4%) in larger firms. For Italy, Contini *et al.* (2009) report that the within-firm standard deviation of wage changes is almost as high as the standard deviation of individual wage changes. Lazear and Shaw (2009b) summarize the findings from this collection of work by emphasizing the striking extent of within-firm variation in wage changes; their overview of the results in Lazear and Shaw (2009a) suggests that the within-firm standard deviation as a proportion of the overall standard deviation of wage changes is of the order of 80–90% in most countries. Sørensen and Vejlin (2011) apply the approach used by Abowd *et al.* (1999) to wage changes using Danish data from 1980 to 2006. They find that worker effects explain almost twice as

much of the variation in wage growth as firm effects. The evidence from these studies suggests that within-firm variation is even more important for earnings changes than for earnings levels.

A recent study by Park and Shin (2017) examines the extent and nature of downward nominal wage rigidity in Korea. They extend this analysis to examining the role of the firm in wage changes, particularly wage cuts. They find that most firms cut wages for some but not all of their employees, indicating strong within-firm heterogeneity of wage changes. This result indicates that the general finding that within-firm effects dominate wage changes applies to cuts as well as rises.

Overall, the evidence seems to suggest that within-firm variation is more important than between-firm variation in explaining the total variation of earnings changes among job stayers. This finding is quite robust across different datasets from different labour markets, including several European countries, the USA and South Korea. Thus despite differences in the institutional setup, different labour markets appear to share a similar reward structure.

In this paper we add to the above literature by examining the role of the firm in explaining heterogeneity in wage changes in Ireland from 2005 to 2013. The fact that this period includes years of both growth and severe recession allows us to examine the stability of the firm's role during a time of macroeconomic instability. Lazear and Shaw (2009b), in their overview of the international evidence, comment that the stability of the within-firm standard deviation of wage changes is an important area for future research. Accordingly, we examine whether and how the breakdown between firm and individual effects is affected by a severe recession.

III. DATA

The analysis is based on data taken from the Job Churn (JC) dataset, which is a longitudinal administrative dataset covering the years 2005–2013 that has been compiled by the Irish Central Statistics Office (CSO). These data combine three elements. Data on annual earnings and weeks worked are provided by the tax authorities in respect of every worker who was an employee during that year. Information on each worker's age, gender and nationality is provided the Department of Social Protection. Finally, data on the sector in which each firm operates and the enterprise's ownership structure come from the CSO's Central Business Register. Anonymized worker and firm identifiers are included in the dataset to allow longitudinal analysis.

There are several significant advantages to using the JC data for our purposes. First, the data cover every employer and almost every employee in Ireland over this period (the top 0.1% of earners were excluded from the data provided to us by the CSO). Since employers are obliged to file tax returns for every worker, problems associated with non-response and attrition are absent, other than as a result of firms going out of business. The fact that the data cover the population of both firms and workers allows us to carry out detailed analyses of within- and between-firm inequalities in Ireland over the period. Moreover, because they are administrative data based on tax returns, the earnings data are largely free from measurement error; it is an offence to misreport workers' earnings in these returns. Finally, the data cover the period from 2005 to 2013, allowing us to compare earnings dynamics before, during and in the immediate aftermath of the crisis.

Earnings are defined in the JC data as annual 'reckonable income' for the calendar year; this is gross income from all sources including bonuses and taxable benefits-in-kind, after non-taxable pension contributions have been deducted. In our analysis, we define

an earnings change as the difference in log earnings between $t - 1$ and t . In keeping with previous work in this literature (e.g. Dickens *et al.* 2007), we define a pay cut as any percentage pay cut greater than 0.1% in absolute value. As with most studies in this area (see, for example, the studies in Lazear and Shaw 2009a), we focus on earnings rather than hourly wages. It might be of additional interest to examine hourly wages, as it is plausible that changes to hours of work occurred in response to the recession. However, the JC data do not contain information on hours of work. Moreover, Doris *et al.* (2015), using an alternative dataset containing hours of work (but no firm identifiers), provide evidence that changes to hours were not the main driver of earnings cuts and pay change heterogeneity. This is in keeping with recent work by Park and Shin (2017) for Korea.

The number of covariates available in the dataset is limited. Information on a worker's gender, age and nationality is provided, as is the industrial sector in which the firm operates. Since we have the population of workers in all firms, we can calculate firm size accurately. However, no data are available on the education level, occupation, labour market experience or job tenure of the workers. To the extent that earnings levels capture human capital characteristics, controlling for initial earnings can be used to overcome the absence of education and labour market experience in our dataset.

Because of the different wage-setting mechanisms that pertain to the public and private sectors, we focus only on private sector workers throughout this paper. In addition, we consider only workers aged 25–60 and only firms with at least 10 such workers. This minimum number of workers is set to ensure that measures of within-firm variation in earnings are meaningful. We use observations on job stayers who work for the same firm for all 52 weeks of both years of a year pair for which a pay change is calculated. Because we have information on only the total number of weeks worked and not when those weeks were worked, restricting the population to those working 52 weeks is the only way we can be certain of excluding seasonal workers from our analysis. In each year, full-year workers comprise about 40% of the total population. Focusing on job stayers ensures that all pay changes can be attributed to a single firm. These restrictions on the subset of workers that we use for the analysis leave us with between 300,000 and 385,000 individual pay change observations and 6000–7000 firms in any given year.

IV. METHODOLOGY

Our first aim is to decompose the heterogeneity of individual pay changes into a part that is due to the firm and a part that is due to the individual's characteristics. We therefore begin with a simple variance decomposition of the dispersion in pay changes into a within-firm component and a between-firm component, following the methodology outlined by Lazear and Shaw (2009b):

$$(1) \quad \sigma_{\Delta w}^2 = \sum_{j=1}^F p_j \sigma_{\Delta w_j}^2 + \sum_{j=1}^F p_j \left(\overline{\Delta w_j} - \overline{\Delta w} \right)^2,$$

where $\sigma_{\Delta w}^2$ is the overall variance of individual earnings changes, $\sigma_{\Delta w_j}^2$ is the variance of earnings changes within firm j , p_j is the share of workers who are working in firm j , $\overline{\Delta w_j}$ is the mean earnings change across workers in firm j , $\overline{\Delta w}$ is the mean earnings change across all workers, and F is the number of firms. Thus the first term of the right-hand side

of equation (1) gives the within-firm variance, which captures the importance of individual characteristics; the second term gives the between-firm variance, which reflects the importance of the firm in determining wage changes. As noted by Song *et al.* (2019) the between-firm component captures all firm-related changes, including changes in the sorting of workers between firms. This variance decomposition is carried out separately for each year pair from 2005/6 to 2012/13. Thus although both the firm and the worker must be observed in both years of a given year pair, firms and workers can enter and exit the population so the same workers and firms are not necessarily observed across all year pairs.

While the decomposition in equation (1) examines pay changes generally, it is also useful to analyse the role of the firm in determining the incidence of pay cuts as discussed earlier. We therefore also employ a decomposition of the variance of pay cuts, as follows:

$$(2) \quad \sigma_c^2 = \sum_{j=1}^F p_j \sigma_{c_j}^2 + \sum_{j=1}^F p_j (\bar{c}_j - \bar{c})^2.$$

In this case, σ_c^2 is the overall variance of a binary variable c that denotes a pay cut, $\sigma_{c_j}^2$ is the variance of pay cuts within firm j , \bar{c}_j is the proportion of workers whose earnings were cut within firm j , and \bar{c} is the proportion of all workers whose earnings were cut; p_j is as defined earlier.

Although these decompositions allow us to examine the relative importance of within-firm and between-firm heterogeneity, they do not explain how individual characteristics affect wage changes within a firm. To examine this issue, we estimate firm fixed effects models of individual wage changes for each pair of years in our sample. In particular, we estimate

$$(3) \quad \Delta w_{ijt} = X_{ij(t-1)} \beta_t + \mu_j + \varepsilon_{ijt},$$

where Δw_{ijt} is the earnings change for worker i in firm j between $t - 1$ and t , and $X_{ij(t-1)}$ is a vector of the characteristics of worker i in firm j at time $t - 1$. The inclusion of the firm fixed effects μ_j means that the estimated β terms capture the average within-firm effect of worker characteristics on individual pay changes. Estimating the μ_j from equation (3) also allows us to re-estimate the proportion of the variance of earnings changes due to firm effects, this time taking into account differences in worker composition across firms.

V. RESULTS

Variance decompositions

As mentioned in Section II, the literature on decomposing wage levels into between- and within-firm parts is more extensive than that on decomposing wage changes. For comparative purposes, we therefore begin our analysis by decomposing wage *levels* into between- and within-firm effects using a decomposition that is analogous to that given in equation (1). The results are given in Table 1. Column (1) indicates how annual earnings inequality evolved in Ireland over the nine-year period. This shows that for full-year workers, inequality was stable from 2005 to 2009, but then began to increase so that by 2013, inequality in annual earnings had increased by just under 10%.

TABLE 1
VARIANCE DECOMPOSITION OF INDIVIDUAL ANNUAL EARNINGS LEVELS

| Year | Overall variance, earnings levels (1) | Between proportion (2) | Within proportion (3) |
|------|--|---------------------------|--------------------------|
| 2005 | 0.387 | 0.42 | 0.58 |
| 2006 | 0.382 | 0.43 | 0.57 |
| 2007 | 0.382 | 0.42 | 0.58 |
| 2008 | 0.383 | 0.42 | 0.58 |
| 2009 | 0.379 | 0.43 | 0.57 |
| 2010 | 0.396 | 0.44 | 0.56 |
| 2011 | 0.412 | 0.46 | 0.54 |
| 2012 | 0.417 | 0.46 | 0.54 |
| 2013 | 0.425 | 0.48 | 0.52 |

TABLE 2
DISTRIBUTION OF INDIVIDUAL ANNUAL EARNINGS CHANGES

| Year | Mean earnings change | S.D. earnings change | Median | 10th percentile | 90th percentile |
|---------|----------------------|----------------------|--------|-----------------|-----------------|
| 2005/6 | 0.068 | 0.137 | 0.058 | -0.042 | 0.196 |
| 2006/7 | 0.068 | 0.140 | 0.068 | -0.046 | 0.199 |
| 2007/8 | 0.048 | 0.144 | 0.045 | -0.069 | 0.175 |
| 2008/9 | -0.022 | 0.155 | -0.006 | -0.171 | 0.103 |
| 2009/10 | 0.003 | 0.144 | 0 | -0.108 | 0.129 |
| 2010/11 | 0.019 | 0.146 | 0.008 | -0.084 | 0.143 |
| 2011/12 | 0.021 | 0.143 | 0.011 | -0.077 | 0.136 |
| 2012/13 | 0.027 | 0.136 | 0.015 | -0.069 | 0.142 |

Columns (2) and (3) of Table 1 show the relative importance of the between-firm and within-firm components of earnings inequality. The results show that prior to the crisis, the relative importance of the firm was stable; in each of these years from 2005 to 2008, firm effects accounted for about 42% of the overall variance. This accords with the research on the variance of wage levels discussed in Section III, which typically indicates that both within-firm and between-firm factors are important, with the within-firm component tending to be slightly stronger.

It is notable, however, that from 2010, the role of the firm became more important; by 2013, the percentage of the growing variation in annual earnings accounted for by the firm was 48%. Thus the growth in annual earnings inequality among job stayers over this period is mainly accounted for by between-firm inequality. However, even during the crisis period, this variance decomposition gives results that are well within the range of those obtained for other countries.

We now turn to the analysis of earnings changes, which is the main focus of the paper. Before carrying out the variance decomposition, we first document the location and spread of the pay change distribution over our data period. Table 2 shows the mean and standard deviation of the difference in log annual earnings for each year pair, as well as the median, 10th percentile and 90th percentile pay changes. The data clearly show the shift to the left of the pay change distribution as a result of the crisis. In 2005/6 and 2006/7, the annual earnings of most job stayers increased substantially, with a mean change of

6.8% in both years. The first signs of the oncoming crisis can be seen in 2007/8 with wage growth slowing, and by 2008/9, both the mean and median wage changes were negative, reflecting the severity of the effects of the recession on workers. Thereafter, some recovery can be seen, but the mean wage change remained below 3% throughout, and the median change was even lower, lying between 1% and 1.5%.

Table 2 also reveals significant heterogeneity in wage growth across individual workers. Even in 2006/7, when the median wage change was 6.8%, the worker at the 10th percentile of the wage change distribution experienced a reduction in earnings of 4.6%, while the worker at the 90th percentile received an earnings increase of 19.9%. Even greater variation is evident at the height of the crisis in 2008/9, with the standard deviation of wage changes increasing to 15.5% from a pre-crisis norm of about 14%. In 2008/9 the median job stayer in Ireland experienced a pay cut of 0.6%. However, workers at the 10th percentile of the wage change distribution received a substantial 17.1% cut in earnings, while those at the 90th percentile saw their earnings increase by 10.3%. While it is evident that many workers experienced significant hardship as a result of the recession, this was not true of every worker.

As noted earlier, one possible explanation for the pattern of pay changes is that the observed heterogeneity reflects differences between firms in pay adjustment strategies. Alternatively, the heterogeneity may reflect differences within firms, as some workers receive large rises while other workers in the same firm receive smaller rises or cuts. This may arise due to differences in worker characteristics that are observed by the firm or due to incentive pay structures that are necessitated by unobserved worker effort. The decomposition presented in Table 3 examines which of these explanations is more important, and clearly comes down in favour of within-firm heterogeneity.

The results show that in the pre-crisis years, the percentage of the variance of earnings changes accounted for by the within-firm component was between 85% and 88%. This figure is very similar to the 87% obtained by Borghans and Kriechel (2009) for the Netherlands. It is also consistent with the conclusion drawn by Contini *et al.* (2009) in the same volume that in most countries, the standard deviation of within-firm pay changes is very similar to the overall standard deviation of pay changes across individuals. It seems that in Ireland, as in other countries, wage changes are linked more to an individual's characteristics than to their firm's fortunes. However, as discussed earlier, to the extent that firms' fortunes interact with worker characteristics, the distinction between the two may be subtle.

The role of individual characteristics may weaken during a downturn if morale issues become more important when pay cuts are under consideration. Given that Ireland was

TABLE 3
VARIANCE DECOMPOSITION OF INDIVIDUAL ANNUAL EARNINGS CHANGES

| Year | Overall variance, earnings change | Between proportion | Within proportion |
|---------|-----------------------------------|--------------------|-------------------|
| 2005/6 | 0.0187 | 0.13 | 0.87 |
| 2006/7 | 0.0196 | 0.12 | 0.88 |
| 2007/8 | 0.0207 | 0.15 | 0.85 |
| 2008/9 | 0.0239 | 0.25 | 0.75 |
| 2009/10 | 0.0207 | 0.18 | 0.82 |
| 2010/11 | 0.0211 | 0.15 | 0.85 |
| 2011/12 | 0.0204 | 0.10 | 0.90 |
| 2012/13 | 0.0185 | 0.10 | 0.90 |

particularly badly hit by the Great Recession, the time period covered by our data provides a useful setting for examining this question. Table 3 shows that with the onset of the crisis, the percentage of the variance of earnings growth accounted for by within-firm factors fell sharply to 75% in 2008/9. To the extent that sectors were affected differentially by the recession, it is plausible that firms reacted differently to the downturn, leading to the firm becoming more important in determining wage changes. This could also have been facilitated by the abandoning of the Irish national wage-setting process in 2009, although as noted earlier, many private sector firms were not adhering to the national agreements for several years prior to the crisis. However, despite the increase in the importance of the between-firm variance in 2008/9, within-firm inequality still dominated the variance of wage changes. Moreover, the drop in the importance of the within-firm variance was very temporary; by 2010/11, it had returned to its previous level of 85%, and then surpassed it in the two subsequent years, rising to 90%. This suggests that the increase in the between-firm variance for 2008–9 reflects a response to the Great Recession rather than any institutional change.

The above decomposition focuses on the variance of wage changes within and between firms but does not pay any attention to the shape of the distribution of these changes. It could be that variances within firms are all similar, resulting in a high proportion of the total variance being attributed to within-firm effects, but that higher moments of the within-firm wage change distribution are quite different. This may particularly arise if negative pay changes are qualitatively different to positive ones; for example, the distribution of earnings changes for firms whose median pay change is negative could be skewed to the left, while for firms giving mostly pay rises, it may be skewed to the right. In this case, the variance of pay changes within firms could be similar, while the shapes of the distributions differ. To examine this issue, we carry out the decomposition of the earnings cut binary variable described by equation (2). This allows us to examine the extent to which the prevalence of wage cuts within a firm explains the heterogeneity in the overall propensity to receive wage cuts.

The results for pay cuts are given in Table 4. The first point to note is that the incidence of pay cuts increased sharply with the onset of the crisis, and then remained high in the following years. Turning to the decomposition itself, we see that the results are in keeping with those presented for pay changes. Although the contribution of between-firm effects increased during the recession, differences within firms still account for over 74% of the variance in wage cuts in 2008/9.

TABLE 4
VARIANCE DECOMPOSITION OF INDIVIDUAL ANNUAL EARNINGS CUTS (BINARY VARIABLE)

| Year | Earnings cuts proportion | Overall variance, earnings cuts | Between proportion | Within proportion |
|---------|--------------------------|---------------------------------|--------------------|-------------------|
| 2005/6 | 0.184 | 0.150 | 0.100 | 0.900 |
| 2006/7 | 0.191 | 0.154 | 0.114 | 0.886 |
| 2007/8 | 0.246 | 0.186 | 0.144 | 0.856 |
| 2008/9 | 0.523 | 0.249 | 0.260 | 0.740 |
| 2009/10 | 0.475 | 0.249 | 0.245 | 0.755 |
| 2010/11 | 0.393 | 0.239 | 0.182 | 0.818 |
| 2011/12 | 0.353 | 0.228 | 0.143 | 0.857 |
| 2012/13 | 0.333 | 0.222 | 0.144 | 0.856 |

To examine the role of the firm in pay cuts in more detail, Figure 1 shows a histogram of the employment-weighted incidence of earnings cuts within firms for 2008/9, the year in which wage cuts peaked. The figure shows that the distribution is almost uniform, with very few firms at or near 0 (no cuts) or 1 (all cuts). Figure 2 shows the same histogram, but without employment weighting. A prominent spike is now visible at 1, which tells us that those firms that did cut pay for everyone tended to be small; perhaps the morale considerations mentioned above are more important within small firms than within large ones. Nevertheless, even in the unweighted histogram, the vast majority of firms—about 90%—cut pay for some workers but not for others. This striking result, which is also found by Park and Shin (2017) using Korean data, does not provide strong support for a morale-based explanation of resistance to pay cuts.

Within-firm regressions

As noted earlier, variation in pay changes within firms can result from heterogeneity in either observable or unobservable characteristics. We now examine the role of available observable factors in explaining wage adjustments within firms, and whether these factors changed during the crisis. To do this, we estimate the fixed effects regression model given in equation (3), regressing individual annual earnings changes on a vector of individual characteristics including age, initial pay (which we may regard as a proxy for initial skill level), gender and nationality. Note that fixed firm-level characteristics cannot be included, as these are not identified once firm effects are introduced. As for the variance decomposition analysis, the observations used are restricted to job stayers aged 25–60 working with the same firm in each year of a given year pair, and to firms with at least 10 such workers; as before, the individuals and firms included can vary across year pairs. The results are given in Table 5.




FIGURE 1. Employment-weighted distribution of within-firm pay cut incidence, 2008/9.

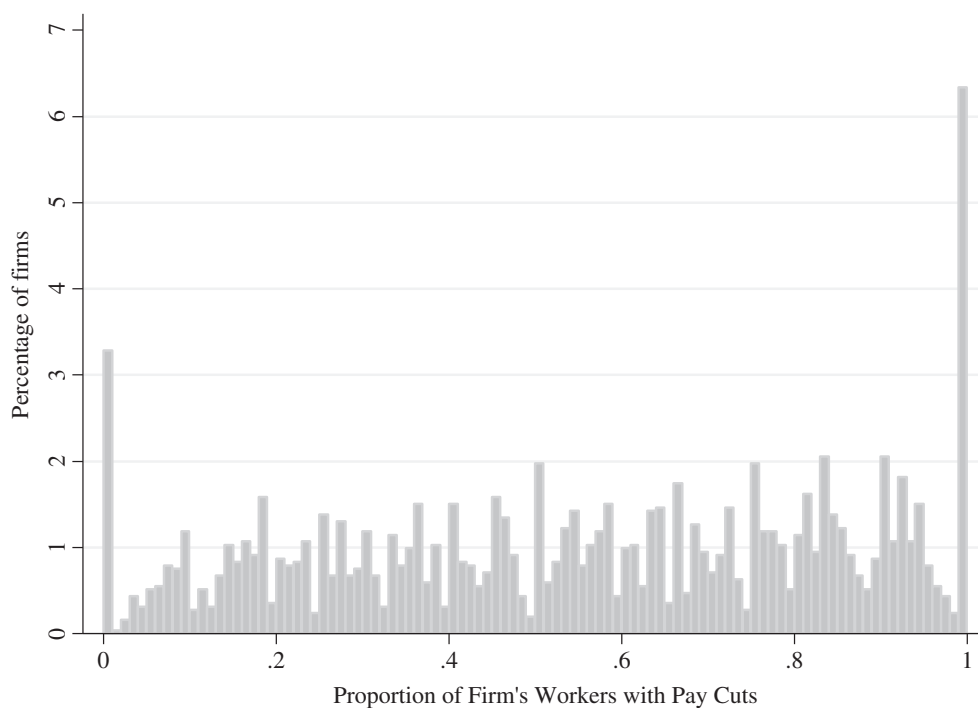


FIGURE 2. Unweighted distribution of within-firm pay cut incidence, 2008/9.

Turning first to the coefficient on initial pay, we see that within firms, pay changes were progressive; higher-paid workers within a given firm received smaller pay rises or bigger pay cuts throughout the period. One explanation for the negative coefficient on initial pay is that firms are concerned with fairness and deliberately protect lower-paid workers from large pay cuts, as suggested by Bewley (1999). However, the estimated coefficient on initial pay will be influenced by the nature of the stochastic component of wages. Random fluctuations in wages from year to year will induce regression to the mean: workers higher up the within-firm earnings distribution in a given year are likely to have arrived there through a combination of talent and good luck. The luck element is unlikely to persist into the following year, resulting in smaller wage rises for better-paid workers and larger pay rises for lower-paid workers. On the other hand, any persistence in the shocks will offset the downward bias associated with regression to the mean. In this context, it is, however, worth noting that in our analysis the estimated coefficient on initial pay becomes markedly more negative in 2008/9. This seems likely to reflect greater employer concerns about fairness as it is not clear why regression to the mean would have become more pronounced during the recession.

With respect to age, the negative coefficients show that older workers receive smaller pay awards within firms in all years, which is consistent with concave age-earnings profiles. However, the size of the coefficient drops substantially in 2008/9. It recovers slightly thereafter, but the coefficient is still half its pre-crisis size by the end of our period. In 2005/6, an employee who was 10 years older than another worker within the same firm would expect to have a pay increase 1.3 percentage points lower than her younger colleague; by 2012/13, her expected pay increase was only 0.7 percentage points lower.

TABLE 5
FIXED EFFECTS MODEL OF INDIVIDUAL ANNUAL EARNINGS CHANGES

| | 2005/6 | 2006/7 | 2007/8 | 2008/9 | 2009/10 | 2010/11 | 2011/12 | 2012/13 |
|-----------------------|---------|---------|---------|---------|---------|---------|---------|---------|
| Initial pay | -0.0303 | -0.0277 | -0.0390 | -0.0578 | -0.0380 | -0.0319 | -0.0390 | -0.0388 |
| Age | -0.0013 | -0.0013 | -0.0011 | -0.0003 | -0.0005 | -0.0006 | -0.0006 | -0.0007 |
| Male | 0.0150 | 0.0151 | 0.0141 | 0.0092 | 0.0122 | 0.0105 | 0.0107 | 0.0109 |
| Irish | -0.0110 | -0.0106 | -0.0060 | -0.0000 | -0.0008 | 0.0013 | 0.0007 | 0.0040 |
| Constant | 0.4334 | 0.4092 | 0.4993 | 0.5973 | 0.4165 | 0.3679 | 0.4465 | 0.4533 |
| Between proportion | 0.201 | 0.171 | 0.212 | 0.336 | 0.233 | 0.232 | 0.161 | 0.157 |

Notes

Given the very large samples used in our analysis, all coefficients are highly statistically significant.

In all years, male workers receive larger pay changes within firms even after controlling for initial pay. This result is consistent with discrimination. However, the result may also reflect males choosing a reward structure with lower pay at the start of their careers in return for higher subsequent earnings growth, with women choosing flatter profiles. Alternatively, it may reflect differences in bargaining power. Card *et al.* (2016) provide evidence that females gain less than their male colleagues when their firm becomes more profitable. It is noteworthy that in our analysis, the male advantage fell from about 1.5 to 1 percentage point with the onset of the crisis, and stayed at this lower level thereafter. It is not obvious why the Great Recession would have induced a reduction in discrimination. On the other hand, it is plausible that the fall in the male coefficient arose because the male advantage in bargaining power during the boom weakened once profits fell and unemployment increased. Since male occupations were particularly hard hit during the recession, differences by gender in employment changes may also be relevant in explaining these findings.

Finally, the effect of Irish nationality on pay changes is interesting. Before the onset of the crisis, being Irish had a negative effect, with Irish workers receiving pay changes about 1 percentage point lower than their non-Irish co-workers. This surprising effect could be due to differences in unobservable characteristics or due to the assimilation of non-nationals over time. However, during the recession, the coefficient became zero, and then turned positive. It could be that the characteristics of Irish workers became more favourable relative to non-Irish workers over the period, perhaps because of differential emigration from the Irish and non-Irish groups.

The last row of Table 5 shows the proportion of the variance due to between-firm effects, after controlling for the observed characteristics. Comparing our results with Table 3, we see that the proportion of the variance attributed to between-firm effects is greater once worker characteristics have been accounted for, as expected. The between-firm variance accounts for between 16% and 34% of the overall variance of pay rises using this methodology, as opposed to 10% to 25% in the unadjusted analysis. However, it is noteworthy that even after controlling for these observed characteristics, the majority of the variation in pay changes is accounted for by within-firm effects. Furthermore, the time patterns in the adjusted and unadjusted variance components are very similar: the between-firm proportion rose substantially in 2008/9, then immediately fell back and was lower at the end of the data period than at the start. Despite the evidence of some changes in personnel practices, the overall impression is that the Great Recession did not permanently alter the role of the firm in wage change determination.

TABLE 6
FIXED EFFECTS MODEL OF INDIVIDUAL ANNUAL EARNINGS CHANGES: FOR ALL FIRMS
AND BY FIRM SIZE

| | 2005/6 | 2006/7 | 2007/8 | 2008/9 | 2009/10 | 2010/11 | 2011/12 | 2012/13 |
|---------------------------|---------|---------|---------|---------|---------|---------|---------|---------|
| <i>Initial pay</i> | | | | | | | | |
| All | -0.0303 | -0.0277 | -0.0390 | -0.0578 | -0.0380 | -0.0319 | -0.0390 | -0.0388 |
| Small | -0.0366 | -0.0357 | -0.0435 | -0.0588 | -0.0415 | -0.0387 | -0.0368 | -0.0370 |
| Medium | -0.0259 | -0.0216 | -0.0328 | -0.0496 | -0.0318 | -0.0240 | -0.0265 | -0.0272 |
| Large | -0.0276 | -0.0247 | -0.0391 | -0.0621 | -0.0392 | -0.0314 | -0.0487 | -0.0477 |
| <i>Age</i> | | | | | | | | |
| All | -0.0013 | -0.0013 | -0.0011 | -0.0003 | -0.0005 | -0.0006 | -0.0006 | -0.0007 |
| Small | -0.0009 | -0.0011 | -0.0010 | -0.0005 | -0.0006 | -0.0005 | -0.0005 | -0.0007 |
| Medium | -0.0014 | -0.0014 | -0.0011 | -0.0004 | -0.0005 | -0.0006 | -0.0007 | -0.0008 |
| Large | -0.0015 | -0.0015 | -0.0012 | -0.0002 | -0.0005 | -0.0006 | -0.0006 | -0.0006 |
| <i>Male</i> | | | | | | | | |
| All | 0.0150 | 0.0151 | 0.0141 | 0.0092 | 0.0122 | 0.0105 | 0.0107 | 0.0109 |
| Small | 0.0134 | 0.0129 | 0.0118 | 0.0055 | 0.0119 | 0.0085 | 0.0103 | 0.0090 |
| Medium | 0.0123 | 0.0147 | 0.0120 | 0.0053 | 0.0098 | 0.0073 | 0.0060 | 0.0081 |
| Large | 0.0178 | 0.0167 | 0.0169 | 0.0139 | 0.0138 | 0.0138 | 0.0142 | 0.0143 |
| <i>Irish</i> | | | | | | | | |
| All | -0.0110 | -0.0106 | -0.0060 | -0.0000 | -0.0008 | 0.0013 | 0.0007 | 0.0040 |
| Small | -0.0097 | -0.0081 | -0.0036 | -0.0005 | -0.0012 | 0.0028 | 0.0031 | 0.0041 |
| Medium | -0.0099 | -0.0113 | -0.0058 | -0.0006 | -0.0001 | 0.0025 | 0.0023 | 0.0051 |
| Large | -0.0127 | -0.0121 | -0.0077 | 0.0006 | -0.0010 | -0.0004 | -0.0024 | 0.0030 |
| <i>Constant</i> | | | | | | | | |
| All | 0.4334 | 0.4092 | 0.4993 | 0.5973 | 0.4165 | 0.3679 | 0.4465 | 0.4533 |
| Small | 0.4845 | 0.4825 | 0.5295 | 0.5881 | 0.4410 | 0.4252 | 0.4093 | 0.4246 |
| Medium | 0.3911 | 0.3461 | 0.4304 | 0.5174 | 0.3557 | 0.2868 | 0.3213 | 0.3374 |
| Large | 0.4148 | 0.3861 | 0.5169 | 0.6522 | 0.4347 | 0.3738 | 0.5577 | 0.5522 |
| <i>Between proportion</i> | | | | | | | | |
| All | 0.201 | 0.171 | 0.212 | 0.336 | 0.233 | 0.232 | 0.161 | 0.157 |
| Small | 0.196 | 0.167 | 0.213 | 0.311 | 0.234 | 0.261 | 0.178 | 0.170 |
| Medium | 0.125 | 0.118 | 0.133 | 0.259 | 0.184 | 0.129 | 0.118 | 0.120 |
| Large | 0.073 | 0.099 | 0.132 | 0.159 | 0.164 | 0.098 | 0.098 | 0.074 |
| <i>Number of firms</i> | | | | | | | | |
| All | 6057 | 6348 | 6763 | 6857 | 7006 | 6856 | 6929 | 7132 |
| Small | 5017 | 5248 | 5575 | 5627 | 5744 | 5631 | 5687 | 5833 |
| Medium | 855 | 912 | 909 | 1002 | 1027 | 997 | 1013 | 1060 |
| Large | 185 | 188 | 219 | 228 | 235 | 228 | 229 | 239 |

Firm size effects

It is not feasible to include firm size as an explanatory variable in the fixed effects model estimated in the previous section; the firm size parameter would be identified from firms who are changing their size substantially and who may not be typical of all firms. However, firm size is potentially an important variable in our analysis, for two reasons. First, firms that are large tend to be foreign-owned and/or exporters to a far greater extent than small and medium-sized firms (Lawless *et al.* 2014); they therefore face different competitive environments than other firms. In addition, larger firms were less likely to face financial distress during the crisis (Lawless and McCann 2013). Second, the wage-change histograms presented in Figures 1 and 2 provide some evidence that across-the-board pay cuts were more common in small firms, perhaps because the morale effects of treating workers differently are stronger in small firms.

We therefore estimate the fixed effects model separately for three firm-size categories: small firms, with 11–49 workers in the first year of the relevant year pair; medium firms, with 50–249 workers; and large firms, with over 250 workers. The results are reported in Table 6. For each variable, the first row replicates the result shown in Table 5; the additional rows are the results of models estimated separately for small, medium and large firms, respectively.

The results indicate some differences in the magnitudes of the coefficients between small, medium and large firms. For example, the initial pay effect is slightly more negative for small firms than for others, while the effect of being Irish is slightly more negative in large firms, and the effect of being male is more positive in large firms. However, overall the time patterns are very similar across all firm sizes.

The last three rows of Table 6 report the proportion of the variance within a firm-size category that is due to between-firm differences, after controlling for observed characteristics. In this respect, it is clear that large firms differ from small and medium firms. Among large firms, an even lower proportion of the total variance is accounted for by between-firm effects than is the case for small or medium firms. This is similar to findings by Borghans and Kriechel (2009), and may indicate that morale considerations are less important in larger firms or that large firms tend to be concentrated in a limited number of industrial sectors and have similar occupational structures.

In summary, although firm identity did become more important during the Great Recession, the overall picture that emerges from the variance decompositions and the fixed effects models is of a labour market in which within-firm forces dominate pay changes.

VI. CONCLUSION

In this paper, we use administrative data to examine the relative role of worker and firm characteristics in explaining the dispersion of wage changes for job stayers in Ireland. Our results show that firm effects play a small role in Ireland, as in other countries for which similar analyses are available. Although firm effects become markedly more important in the peak year of the economic crisis, the vast majority of earnings changes continue to be driven by individual worker characteristics. This finding runs contrary to our expectations. When we examine this further by investigating the incidence of pay cuts directly, we find a stronger role for the firm, but nevertheless, the vast majority of pay cuts occur in firms where other workers' pay is not cut.

To control for differences in worker characteristics within firms, we estimate a firm fixed effects model of individual pay changes. We find evidence of changing personnel practices as a result of the Great Recession. These changes occur across all firm sizes and in many cases persist beyond the immediate onset of the crisis. When account is taken of worker characteristics, the role of the firm in determining pay changes increases. But it remains the case that most of the variance of individual pay changes is accounted for by worker characteristics rather than the identity of the firm.

The prevalence of wage cuts in Ireland during the Great Recession offers a rare opportunity to examine the behaviour of firms in a distressed labour market. Morale considerations lead us to expect that in such circumstances, firms that cut wages will do so for all or most workers, so that between-firm effects would be the dominant component of the variance of pay cuts. While we did see an increase in the between-firm component with the onset of the crisis, within-firm effects still explained over 74% of the variation in wage cuts. This suggests that morale issues may be less important than

employers reported to Bewley (1999), perhaps because workers within firms are segregated along occupational lines across which pay comparisons are unlikely. Alternatively, an individual worker's morale may be less negatively affected by a pay cut when there are numerous reports of pay cuts in the wider economy, even if the pay of other co-workers is not being cut. Although employment changes are not considered in this paper, perhaps the negative morale implications of workers within the same firm being given differential pay changes are entirely offset by workers' gratitude for retaining their jobs at a time of rapidly increasing unemployment.

ACKNOWLEDGMENTS

We would like to thank Ignacio Garcia Perez, seminar participants at UCD, the New Zealand Work Research Institute, the 2016 Irish Economic Association Annual Conference and two anonymous referees for helpful comments on an earlier draft of this paper.

REFERENCES

- ABOWD, J. M., KRAMARZ, F. and MARGOLIS, D. N. (1999). High wage workers and high wage firms. *Econometrica*, **67**, 251–333.
- ALVAREZ, J., BENGUARIA, F., ENGBORN, N. and MOSER, C. (2018). Firms and the decline in earnings inequality in Brazil. *American Economic Journal: Macroeconomics*, **10**, 149–89.
- BAKER, G., GIBBS, M. and HOLMSTROM, B. (1994). The wage policy of a firm. *Quarterly Journal of Economics*, **109**, 921–55.
- BARTH, E., BRYSON, A., DAVIS, J. and FREEMAN, R. (2016). It's where you work: increases in earnings dispersion across establishments and individuals in the U.S. *Journal of Labour Economics*, **34**, S67–S97.
- BEWLEY, T. F. (1999). *Why Wages Don't Fall During a Recession*. Cambridge, MA: Harvard University Press.
- BORGHANS, L. and KRIEHEL, B. (2009). Wage structure and labor mobility in the Netherlands, 1999–2003. In E. P. LAZEAR and K. L. SHAW (eds), *The Structure of Wages: An International Comparison*. Chicago, IL: University of Chicago Press, pp. 125–48.
- CARD, D., HEINING, J. and KLINE, P. (2013). Workplace heterogeneity and the rise of West German wage inequality. *Quarterly Journal of Economics*, **128**, 967–1015.
- , CARDOSO, A. and KLINE, P. (2016). Bargaining, sorting and the gender wage gap: quantifying the impact of firms on the relative pay of women. *Quarterly Journal of Economics*, **131**, 633–86.
- CARDOSO, A. R. (1997). Workers or employers: who is shaping wage inequality? *Oxford Bulletin of Economics and Statistics*, **59**, 523–47.
- (1999). Firms' wage policies and the rise in labor market inequality: the case of Portugal. *Industrial and Labor Relations Review*, **53**, 87–102.
- (2000). Wage differentials across firms: an application of multilevel modelling. *Journal of Applied Econometrics*, **15**, 343–54.
- CONTINI, B., LEOMBRUNI, R., PACELLI, L. and VILLOSIO, C. (2009). Wage mobility and dynamics in Italy in the 1990s. In E. P. LAZEAR, and K. L. SHAW (eds), *The Structure of Wages: An International Comparison*. Chicago, IL: University of Chicago Press, pp. 373–99.
- DICKENS, W. T., GOETTE, L., GROSHEN, E. L., HOLDEN, S., MESSINA, J., SCHWEITZER, M. E., TURUNEN, J. and WARD, M. E. (2007). How wages change: micro evidence from the international wage flexibility project. *Journal of Economic Perspectives*, **21**, 195–214.
- DORIS, A., O'NEILL, D. and SWEETMAN, O. (2015). Wage flexibility and the Great Recession: the response of the Irish labour market. *IZA Journal of European Labor Studies*, **4**(18).
- DUSTMANN, C. and MEGHIR, C. (2005). Wages, experience and seniority. *Review of Economic Studies*, **72**, 77–108.
- ELSBY, M. W. L., SHIN, D. and SOLON, G. (2016). Wage adjustment in the Great Recession and other downturns: evidence from the United States and Great Britain. *Journal of Labor Economics*, **34**, S249–S291.
- LAWLESS, M. and MCCANN, F. (2013). Determinants of default: evidence from a sectoral-level panel of Irish SME loans. *Economic and Social Review*, **44**, 473–88.
- , ——— and MCINDOE-CALDER, T. (2014). SMEs in Ireland: stylised facts from the real economy and credit market. *Policy Studies*, **35**, 435–57.
- LAZEAR, E. P. and SHAW, K. L. (2009a). *The Structure of Wages: An International Comparison*. Chicago, IL: University of Chicago Press.

- and ——— (2009b). Wage structure, raises and mobility: an introduction to international comparisons of the structure of wages within and across firms. In E. P. LAZEAR and K. L. SHAW (eds), *The Structure of Wages: An International Comparison*. Chicago, IL: University of Chicago Press, pp. 1–57.
- PARK, S. and SHIN, D. (2017). The extent and nature of downward nominal wage flexibility: an analysis of longitudinal worker/establishment data from Korea. *Labour Economics*, **48**, 67–86.
- SONG, J., PRICE, D. J., GUVENEN, F., BLOOM, N. and von WACHTER, T. (2019). Firming up inequality. *Quarterly Journal of Economics*, **134**(1), 1–50.
- SØRENSEN, K. L. and VEJLIN, R. (2011). Worker and firm heterogeneity in wage growth: an AKM approach. *LABOUR*, **25**, 485–507.