Who Pays for Occupational Pensions?*

Ola Lotherington Vestad†

Ragnar Frisch Centre for Economic Research

12th March, 2012

Abstract

The purpose of this paper is to estimate wage effects of occupational pensions, exploiting the introduction of mandatory occupational pensions in Norway as a source of exogenous variation in pension coverage. Various difference-in-differences models are estimated on a large sample of Norwegian private sector firms. The results indicate that on average, less than half the costs of a minimum requirement occupational pension was shifted from firms to workers in terms of lower wages, and that there are important heterogeneities with respect to the influence of local unions and central negotiations on the wage setting in different industries.

JEL codes: H22, J32, J38, J50.

Keywords: pension reform, mandatory occupational pensions, labour unions and centralised negotiations, matched employer-employee register data.

*This paper is part of the Frisch Centre project *1307 Strategic research programme on pensions* which is financed by the Ministry of Labour. Data made available to the Frisch Centre by Statistics Norway has been essential for the project. I am grateful to James Banks, Bernt Bratsberg, Harald Dale-Olsen, Steinar Holden, Erik Hernæs, Oddbjørn Raaum, Knut Røed, Torgeir Aarvaag Stokke and Kjetil Storesletten for valuable comments and guidance. Comments from seminar participants at the Norwegian Economist Meeting 2011, University of Oslo, Institute for Fiscal Studies, Norwegian School of Management, the 25^{th} Annual ESPE Conference, the 26^{th} Annual EEA-ESEM Conference, the 13^{th} IZA/CEPR European Summer Symposium in Labour Economics and the 23^{rd} Annual EALE Conference are also gratefully acknowledged. I also wish to thank the Institute for Fiscal Studies for its great hospitality during a stay in the spring of 2011. The usual disclaimer applies.

†ola.vestad@frisch.uio.no

1 Introduction

The purpose of this paper is to estimate the wage effects of occupational pensions (OPs, or employer-provided pensions) for a large sample of Norwegian private sector firms. Knowledge about the offset between pensions and wages is becoming increasingly important, as OPs are expected to play a more prominent role in retirement provision in many countries where governments seek to reduce their pension commitments. In Norway, mandatory OPs were introduced in 2006 as part of an ongoing reform of the public pension system. One of the arguments behind this new mandate was based on a concern that workers not covered by an OP would be left with insufficient pension benefits under a new and less generous public pension regime. Mandatory OPs along with cuts in public pensions may be a politically attractive alternative to higher taxes and less drastic cuts in public pensions, and an assessment of the extent to which the costs of such mandated benefits are shared between firms and workers is crucial if one wants to know who bears the costs of pension reforms.

Occupational pensions are part of compensation packages offered by firms to workers, and identification of the offset between pensions and wages is complicated by the joint determination of pensions, wages and other forms of compensation. The positive coefficient on pensions which is typically found in cross section wage regressions (see e.g. Hernæs et al. (2010)) is thus likely to be corrupted by simultaneity bias and the imperfect observability of productivity. In this paper, a difference-in-differences strategy is used to exploit the introduction of mandatory OPs in Norway as a source of exogenous variation in pension coverage. Provided that the counter-factual trends in firm level wages are independent of pre-reform OP status, conditional on observed covariates, this approach gives unbiased and consistent estimates of the offset factor between pensions and wages.

As only half the workers in the Norwegian private sector were covered by an occupational pension prior to the reform, a natural question to ask is the following: What motivates some firms to offer an OP while others choose not to? Gustman, Mitchell, and Steinmeier (1994) point to several possible reasons, one being that firms offer OPs simply because they are demanded by workers. Such a demand from workers may be motivated by the fact that both contributions and benefits tend to be tax-favoured, so that the after-tax return to savings in OP schemes may exceed the return earned in other savings vehicles. This feature makes OPs more attractive to high-wage workers than to workers with lower wages, given that taxes on wage income are progressive. There may also exist economies of scale, making group-saving more cost effective than individual saving, which may help explain the stylized fact that OPs have been a large firm phenomenon. A third reason why workers may desire OPs is that they often provide insurance of a type that is hard to obtain otherwise than through OP schemes, such as for instance disability insurance.

Occupational pensions may also be used by firms as a means of minimizing labour costs and increasing productivity. In a setting with lifetime contracts, Lazear (1981) argued that deferred compensation could be used to minimize the cost of inducing optimal effort in firms where monitoring effort is difficult or costly. In the absence of lifetime contracts, firms faced with substantial hiring and training costs could use OPs to discourage turnover and/or as a means of attracting stayers rather than movers. Defined benefit (DB) pensions are particularly well suited for this purpose, as they are designed to give a certain proportion of the final wage as yearly pension benefits, and as DB pension covered workers typically face capital losses if they leave the firm prior to retirement (see Hernæs et al. (2011)). Finally, OPs may be used to induce retirement for workers whose productivity falls by age, by making retirement more economically attractive.

Hernæs et al. (2010) provide evidence that supports several of the hypotheses above: Occupational pensions are typically found in large firms, in firms where tax gains to the employees are high, and in firms where long periods of training are required. They also find that the occurrence of an OP increases tenure substantially, and that firms with the most to gain in terms of higher expected tenure were those who actually chose to have an OP. The data used for the analyses in this paper does not contain direct information on all the above-mentioned aspects, but the available set of covariates should nevertheless be sufficient to account for the main differences between firms with and without OPs prior to the reform.

As for the economic incidence of mandatory OPs, the question asked in this paper is whether or not it falls on workers in terms of reduced wages. Employers are obliged by law to cover the direct costs, consisting of contributions, waiver of contribution (to cover continued contributions in the event of disability) and administrative costs, but they may well have been able to shift parts of these costs onto employees. If the full costs are borne by the affected workers one would expect to see no effects on employment, and one would not need to be concerned about redistributive effects of the mandate. Given that the estimated average treatment effects in this paper indicate that only half the costs of a minimum requirement OP is passed on to workers in terms of lower wages, one could ask whether there have been adjustments in non-wage amenities other than pensions. There is no detailed information on non-wage amenities in our data, but we note that crude measures of the prevalence of non-wage amenities provided by Statistics Norway (http://www.ssb.no/english/) show that the fraction of workers receiving such types of compensation has gone up rather than down over the relevant period.

Firms that are unable to shift the full costs onto workers by adjusting compensation could either adjust employment, or try to pass it on to consumers in terms of price adjustments or to firm owners through reduced profits. In each of these cases, the mandate may have redistributive effects. One might see redistribution from less productive to more productive workers, if low productivity workers are forced to reduce their hours of work, or from consumers or firm owners to previously uncovered workers. This paper, however, is devoted exclusively to an investigation of the wage effects of the mandate.

The paper is organized as follows: Section 2 reviews some related theoretical and empirical literature, before a brief history of occupational pensions in Norway is given in Section 3, including a description of the process that materialised into mandatory OPs. This background information is essential to the interpretation of the empirical results of the analysis. Section 4 describes the three sources of data that are used to create a sample consisting of 10,392 Norwegian private sector firms, and gives some descriptive statistics. Firms with no OP in place prior to the reform constitute 40% of the firms in the sample, and these differ from the OP-firms in a number of ways: They are smaller, pay lower wages and employ younger and less educated workers.

The empirical specifications are spelled out in Section 5, and the main focus from there on is on various difference-in-differences models. Estimation results given in Section 6 indicate that there has been some cost sharing between firms and workers, but the costs of occupational pensions are not fully shifted onto workers. The fixed effects estimates are reduced by about 60% when a small number of firms with very high or very low estimated propensities to have an OP prior to the reform are removed from the sample. I argue that this difference is mostly due to very high wage growth among the firms with the highest estimated propensities, and base the further investigations described in Section 7 on a reduced sample consisting of more comparable firms. Results from a specification that allows for both pre- and post-reform effects indicate that firms were holding back on wages several years before the reform was formally implemented, but the closest we get to a full shift is about 50% of costs shifted onto workers (in 2009). The final specification takes into account these dynamic adjustments at the same time as it allows the treatment effects to vary with measures of the share of unionised workers and the influence of central negotiations on wages in different activities. More appears to be shifted onto workers in activities where wages are influenced by central, but not local negotiations, and in activities with low shares of unionised workers.

Section 8 concludes and gives some prospects for further work.

2 Related literature

Previous studies of the pension-wage offset have often been based on a compensating differentials framework; see e.g. Montgomery, Shaw, and Benedict (1992), Smith (1981) and Schiller and Weiss (1980). Although the hypothesis of a one-to-one compensating differential between pensions and wages could not always be rejected, the evidence in favour of the theory is not overwhelming. These studies were bound to rely on rather small and unrepresentative cross sectional samples of workers, and such data limitations make credible identification difficult, if not impossible, given the joint determination of pensions and wages and the imperfect observability of worker productivity.

As for the question of who ends up bearing the costs of mandated benefits, the existing economic literature provides no single clear cut answer. Assuming first that the mandated benefit is perceived as a regular tax both by workers and by firms, that is, that workers assign no value to the future pension benefits, we know from the textbook example of tax incidence in competitive markets that the "least elastic" side of the market ends up paying most of the tax. Also, if the labour supply curve for some exogenous reason becomes steeper ("less elastic"), the model predicts that a bigger proportion of the tax burden will be shifted onto workers through lower wages, and effects on employment will be lower. Summers (1989) pointed out that if workers assign some value to the mandated benefit, one may expect a positive shift in labour supply when the mandate is implemented. Again, more of the costs would be shifted onto workers and effects on employment would be lower, along with a reduced dead weight loss.

Summers, Gruber, and Vergara (1993) argued that unions are more likely than individuals to recognise the link between contributions paid and benefits received, implying that a shift in labour supply is more likely to occur in labour markets in which unions play a central role, as they do in countries like Norway. In the limiting case where workers' (or unions') valuation of the mandated benefit is the same as its cost to firms, the entire cost will be shifted onto workers, there will be no effects on employment and thus no dead weight loss associated with the mandate. Mandating OPs would then be a more efficient way of securing sufficient pensions for otherwise non-covered workers than publicly provided pensions financed through taxes. Another part of this story is that if the costs of OPs are not fully shifted onto workers, it must be either because OPs are not sufficiently valued by workers or unions, or because there are impediments to the adjustment of relative wages to reflect workers' valuation.

Alesina and Perotti (1997) present a theoretical model that predicts a hump-shaped relationship between the degree of centralisation, defined as the inverse of the number of unions in the economy, and the degree of shifting of labour taxation. The intuition behind this relationship is as follows: Larger parts of the labour tax burden is borne by employees in competitive labour markets, with inelastic individual labour supply, and in centralised economies, where a small number of unions internalise macroeconomic constraints and effects of wage increases on labour costs and employment, than in economies with intermediate levels of centralisation, where unions are large enough to have significant impacts on wages but at the same time too small or too numerous to properly internalise the adverse effects of bargaining outcomes. The theoretical predictions are supported by empirical evidence from data on the manufacturing sector in 14 OECD countries.

Turning to the empirical literature on the incidence of mandated benefits, a notable example is Gruber (1994), who studied the economic incidence of mandated maternity benefits in the US. His findings consistently suggest full shifting of the costs of the mandates and he found little effect on total labour input for the groups of workers affected by the mandates. A more recent contribution is a study of the incidence of social security contributions by Ooghe, Schokkaert, and Flechet (2003), based on sectoral panel data covering six different European countries. Testing the predictions of an efficient bargaining model they find that at least 50% of both legal and customary contributions are shifted onto workers, and suggest that these results are due to trade unions recognizing the link between contributions and benefits during wage negotiations. Finally, a study of the incidence of non-wage labour costs in OECD countries by Azémar and Desbordes (2010) establishes that in countries with highly coordinated bargaining, the entire tax burden appears to be shifted immediately onto workers. We will return to the impacts of centralised negotiations and local labour unions in Section 7, after having established the back bones of the empirical framework.

3 Institutional background

The Norwegian pension system may be seen as one consisting of three different layers, of which the National Insurance Scheme (NIS) constitutes the first. The NIS provides universal coverage, meaning that all citizens above the age of 67 are guaranteed a minimum pension. On top of this comes an earnings related pension for those who have had sufficient earnings throughout their career. The second and third layers are occupational pension schemes and voluntary individual savings.

As for the occupational pensions there are separate systems for the public and the private sectors. The public sector scheme is of the defined benefit[1] type, and guarantees yearly pension payments corresponding to 66% of the final yearly wage income after 30 years of service. The market for OPs in the private sector used to be strongly dominated by DB pensions, as contributions qualified as a tax deductible cost only for pension plans of this type until 2001, when there was a change in legislation. Since then, firms' contributions to both defined benefit and defined contribution plans are treated like wages for tax purposes, provided the plan meets a set of requirements imposed by the government. For employees, both contributions and accumulations are tax exempted, while benefits are taxed under the income tax, but at a lower rate than wage income.

The new legislation appears to have led to a marked increase in the number

[1] DB pensions schemes are designed such that they guarantee or target a certain level of pension benefits, defined as a proportion of final yearly wage income. In Defined Contribution (DC) pension schemes there is no such guarantee or target; they specify in stead the annual contributions as a proportion of wages.

of firms operating a DC plan, but the majority of firms establishing DC plans were already operating a DB plan, which would typically be closed for new entrants as soon as the new DC plan was in place. The increase in individual OP coverage and in the number of firms operating OPs followed by the new legislation was thus fairly moderate (see e.g. Midtsundstad and Hippe (2005) and Veland (2008)). Motivated by a concern for non-covered workers, the labour unions made a first proper attempt to establish a collective OP scheme for the private sector during the central negotiations in 2002. This did not succeed, but a settlement on mandatory OPs was part of the outcome of the central negotiations two years later, in 2004. The final result of this settlement was the Act relating to mandatory occupational pensions, which entered into force on January 1 2006.

While only about 50% of workers in the private sector were covered by an OP prior to 2006, the Act required all firms except small businesses, self employed and family businesses to have an OP in place by the end of 2006, with economic effect for the employees as of July 1 2006. All employees working at least 20% of full time are granted membership in the firm's pension plan, and the Act also specifies a minimum contribution of 2% of earnings between 1 and 12 Basic amounts (G)[2]. It appears that most of the firms that were forced by the Act to introduce an OP chose the minimum level of generousity (see Veland (2008)). Another central element of the Act is that employers are obliged to cover the direct costs related to the OP, including contributions, waiver of contribution (to cover continued contributions in the event of disability) and administrative costs. For a minimum requirement OP these costs amount to about 2.6% of earnings.

[2]The Basic amount is frequently referred to as G, and is a central feature of the public pension system in Norway. G is adjusted every year, with a nominal rate of growth varying between 2 and 14% since its introduction in 1967. The average G for 2010 was 74721 NOK, which corresponds to about 9300 EUR or 8000 GBP. For further details on G and on the public pension system in general, see e.g. Iskhakov (2008).

4 Data, sample and descriptive statistics

4.1 The data

The empirical analysis is based upon three sources of data; the Register of Employers and Employees (REE), and pension liabilities and costs, respectively, from a set of balance sheet data and a set of accounting data. The REE is a linked employer-employee data set based on administrative registers, and covers the entire Norwegian working-age population over the period 1992-2009. For each pair (employer, employee), the REE contains a wide range of both individual and firm specific information, such as age, hours worked, earned income, opening and closing dates for the employment record, industry code, geographical location and organizational structure.

The second source of data contains enterprise based financial information recorded by the authorities, and covers the period 1999-2005. Enterprises with defined benefit occupational pension plans have to set aside assets to cover pension liabilities. By the end of each year, when the annual accounts for the enterprise are made up, pension assets and liabilities are usually not identical, and over- or under-funding will enter the balance. This is what is used for identification of enterprises operating a DB pension plan.

A third source of data containing enterprise level accounting data is used to identify firms with a defined contribution pension plan. These data cover the years 2005 through 2008. DC pension costs are treated much the same way as regular wages, and thus do not appear in the balance. Firms having pension costs registered in the profit-and-loss account and no entries for over- or under-funding of DB plans in the balance are identified as DC firms.

4.2 The sample

Based upon the union of the three data sources described above I define a panel of firms covering the period 1999-2009. Firms are included in the panel only if they satisfy the following criteria: (i) They were in operation in 2005, and (ii) they had at least ten full-time, full-year employees[3], each earning more than 100,000 NOK. Criterion (i) is necessary for classification of firms with respect to their OP status, as 2005 is the only year for which information is available from all of the three data sources.[4] The second criterion is imposed to make sure that all firms in the panel were covered by the Act, and must be satisfied for each pair (firm, year). Setting the minimum number of employees as high as at ten leaves out pure family businesses not covered by the Act, and newborn firms that simply did not have the time to establish an occupational pension prior to 2005.

We restrict attention to the private sector, and divide the private sector firms into three separate groups according to their OP status by the end of the base year (2005): DB-firms, DC-firms and No OP-firms. The panel consists of 10,392 firms, and among these were 40 percent without an OP prior to the Act, while 27 percent had a DC plan and about 32 percent of the firms had a private sector DB pension plan (Table 1). 62 percent of the employees associated with these firms were covered by a private sector DB pension, which confirms the stylized fact that occupational pensions of the DB type is a large firm phenomenon. The fractions of firms and employees in the four sectors remained fairly stable across the observation period.

	Firms		**Employees**	
OP status	n	Per cent	n	Per cent
DB	3,394	32.66	287,956	62.22
DC	2,836	27.29	98,159	21.21
No OP	4,162	40.05	76,691	16.57
All	10,392	100.00	462,806	100.00

Table 1: The number of firms and employees by OP status (2005)

More detailed descriptive statistics are given in Table 2. Firms without an OP in 2005 are clearly different from the two other groups of firms in several

[3]Employees younger than 20 or older than 66 years are not counted.

[4]More precisely, the first criterion is that the firms were present in all three data sets in 2005. Relaxing this requirement and counting firms with unknown OP status as "No OP-firms" increases the sample size by about 800 firms, but most likely also the risk of misclassification.

aspects: They are much smaller, they pay lower wages, they have even higher fractions of male employees, their employees are younger, have less tenure and are less educated, firms operating in construction are heavily over-represented, and relatively few of them are based in Oslo (the capital).

Table 2: Summary statistics - firm level means by OP status (2005)

Variable	**DB**	**DC**	**No OP**
No. of employees[a]	84.843	34.612	18.426
ln(Mean earnings)[b]	12.848	12.821	12.736
Fraction of males	0.729	0.714	0.772
Avg. age of employees	43.306	42.048	39.967
Avg. tenure	7.767	6.847	6.148
Fraction of highly educated	0.223	0.273	0.164
Fraction of immigrants	0.055	0.056	0.068
No. of years in panel	9.312	8.490	7.588
Selected industry dummies			
Manufacturing	0.336	0.199	0.192
Construction	0.080	0.138	0.231
Wholesale and retail trade, ...	0.250	0.227	0.252
Real estate, renting and business activities	0.103	0.187	0.148
Selected county dummies			
Akershus	0.118	0.095	0.097
Oslo	0.224	0.224	0.169
Rogaland	0.087	0.088	0.093
Hordaland	0.082	0.098	0.092
Number of firms	3,394	2,836	4,162

[a] Full-time, full-year employees.
[b] Earnings of full-time, full-year employees, measured in terms of 2004-prices.

Figure 1 shows box plots[5] of firm level average wages by OP status (in 2005) and for each of the eight years. Group level averages of log-wages are plotted in Figure 2. The two figures show that there has been some growth in real wages in all three groups over the period of observation. Also, the spread of real wages appear to have increased, and the observed (unadjusted) trends in log-wages do not seem to be too different across OP and No OP firms.

[5]The lower and upper hinges of the boxes indicate the 25^{th} and 75^{th} percentiles, respectively, denoted by $x_{[25]}$ and $x_{[75]}$, and the horizontal lines cutting through the boxes indicate the median. The vertical lines below and above the boxes are called adjacent lines, and the markers on each end of the lines indicate lower and upper adjacent value, respectively. Adjacent values are calculated as described in the Stata Manual [G] Graphics: Define x_i as the ith ordered value of x, and define U as $x_{[75]} + \frac{2}{3}(x_{[75]} - x_{[25]})$ and L as $x_{[25]} - \frac{2}{3}(x_{[75]} - x_{[25]})$. The upper adjacent value is x_i such that $x_i \leq U$ and $x_{i+1} > U$, and the lower adjacent value is x_i such that $x_i \geq L$ and $x_{i+1} < L$. Observations above (below) the upper (lower) adjacent values are not shown in the figure.

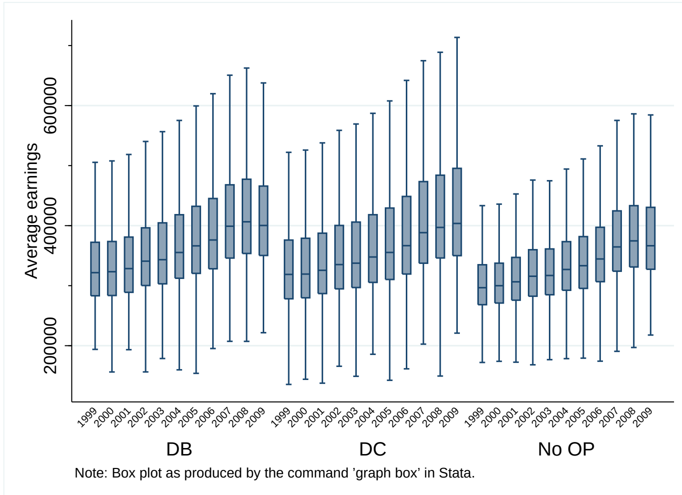

Figure 1: Firm level average earnings by year and OP status (2005). Earnings measured in terms of 2004-prices.

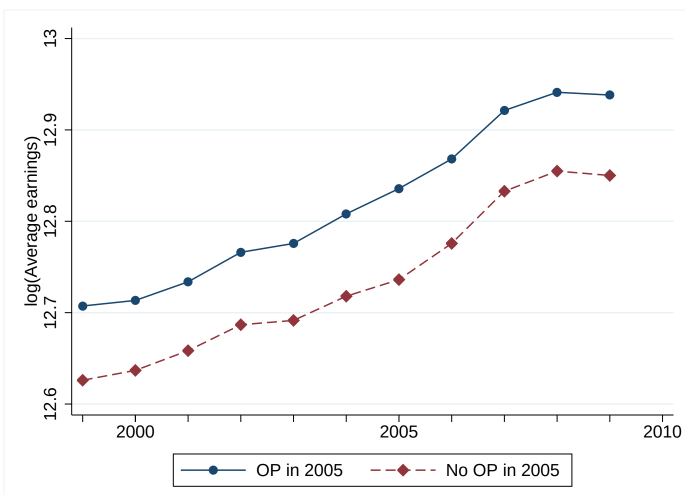

Figure 2: Log of firm level average earnings by year and OP status (2005). Averages across firms within OP group, earnings measured in terms of 2004-prices.

5 Empirical specification and identification

The identification strategy is built upon the idea of exploiting the introduction of mandatory occupational pensions as a source of exogenous variation in pension coverage. A difference-in-differences way of thinking seems promising in this setting. Let $\bar{w}_{jst}$ denote the outcome of interest, specified as the natural log of the average wage among full-time, full-year employees in firm j in group s at time t.[6] The observed wage is $\bar{w}^0_{jst}$ for the non-treated and $\bar{w}^1_{jst}$ for the treated, and only one of these is observed for each firm. The key identifying assumption is that the counter factual trend behavior of log-wages is the same for the treatment and control groups, conditional on observed covariates.[7] In other words, we assume that the growth rate in log-wages before and after the reform would have been the same for No OP-firms as for OP-firms, in the absence of reform. Now, as treatment status varies only at the group level in this particular case, one may argue that the source of omitted variable bias is most likely to be unobserved variables at the OP-group and year level. The idea behind the difference-in-differences identification strategy is that these group-level omitted variables can be captured by group-level fixed effects.

Assume that in the absence of reform, the outcome variable is determined by a time-invariant group or OP-status effect (γ_s), a common year effect (λ_t) and observed firm-specific covariates (X_{jt}), that is,

$$E\left[\bar{w}^0_{jst}|s,t,X_{jt}\right] = \gamma_s + \lambda_t + X^{'}_{jt}\beta$$

As a first approach, we assume that the effect of treatment is additive and

[6] Both the notation and the framing of the estimation strategy in this section is borrowed from Angrist and Pischke (2009).

[7] We also need two other assumptions to be satisfied: First, the reform must be exogenous in the sense that it did not just institutionalise a pre-existing trend towards broader OP-coverage. This assumption is likely to hold true, cf. the discussion in Section 3. Second, there can be no spill-over or general equilibrium effects of the reform. This assumption would be questionable if for instance the reform led many firms to close down, which in turn would lead to increased supply of certain groups of workers and therefore to lower wages for the same groups of workers and their substitutes.

constant, denoted by δ, so that

$$E\left[\bar{w}^1_{jst}|s,t,X_{jt}\right] = E\left[\bar{w}^0_{jst}|s,t,X_{jt}\right] + \delta$$

Together these assumptions imply that the observed outcome may be written as

$$\bar{w}_{jst} = \gamma_s + \lambda_t + \delta D_{st} + X^{'}_{jt}\beta + \varepsilon_{jst}, \tag{1}$$

where $E\left[\varepsilon_{jst}|s,t,X_{jt}\right] = 0$ and D_{st} is an indicator for treatment status, the regressor of interest. The population difference-in-differences is

$$\begin{aligned}\delta = &\{E\left[\bar{w}_{jst'}|s = NoOP, t', X_{jt'}\right] - E\left[\bar{w}_{jst}|s = NoOP, t, X_{jt}\right]\} \\ &- \{E\left[\bar{w}_{jst'}|s = OP, t', X_{jt'}\right] - E\left[\bar{w}_{jst}|s = OP, t, X_{jt}\right]\},\end{aligned}$$

where t and t' denote before- and after-reform observations. δ thus gives us the difference in log-wages before and after the reform for those directly affected by the reform (the treatment group) relatively to those who were not affected by the reform (the control groups), taking account for observed firm-specific characteristics, unobserved group fixed effects and economy-wide factors potentially affecting the various groups over time.

It may well be that the 'group level fixed effects' way of thinking is too rough, i.e. that unobserved factors at the firm level are of great importance for the wage levels. Assuming that these factors are constant over time, and retaining the assumptions of equal counterfactual trend behavior in the treatment and control groups and of constant and additive treatment effect, we specify a modified version of (1) as follows:

$$\begin{aligned}&\bar{w}_{jt} = \alpha_j + \lambda_t + \delta D_{jt} + X^{'}_{jt}\beta + \varepsilon_{jt}, \text{ where} \\ &\varepsilon_{jt} = \bar{w}^0_{jt} - E\left[\bar{w}^0_{jt}|F_j, t, X_{jt}\right], \text{ and} \\ &\alpha_j = \alpha + F^{'}_j\rho\end{aligned} \tag{2}$$

F_j are the unobserved firm level fixed effects, and X_{jt} are observed time varying covariates.

6 Estimation

Let I_s be an indicator for firms who did not have an occupational pension plan in 2005, and d_t be a dummy taking the value 1 for observations after the reform, and 0 otherwise. The first equation to be estimated is a version of equation (1):

$$\bar{w}_{jkst} = \alpha + \gamma_s \mathit{OP\text{-}group}_s + \lambda_{kt} t + \delta \left(I_s \cdot d_t \right) + X_{jt}^{'} \beta + \varepsilon_{jkst}, \tag{3}$$

where $I_s \cdot d_t = D_{st}$ and λ_{kt} is an industry-specific trend coefficient multiplying the time-trend variable t. The vector X_{jt} contains a cubic term in the mean age of the employees (as a proxy for the general level of experience), the fraction of male employees, the fraction of highly educated employees (where a high level of education is defined as education at the bachelor level or above), the fraction of immigrants (first and second generation) and dummies for geographical location (county).

(3) is more general than (1), in that it allows for industry specific year effects (or flexible time trends). A rationale for including industry-specific trends rather than just a common trend is the 2004 expansion of the European Union, an event with the potential of having different effects on different industries. Bratsberg and Raaum (2010) show how this expansion represented a positive shift in aggregate supply of immigrant workers to the Norwegian construction industry, leading to lower wage growth for workers in trades with rising immigrant employment than for other workers, and an outflow of low wage workers from the industry. Although 23% of the NoOP-firms in the panel used in this study are associated with the construction industry, switching from a common trend to industry specific trends in a specification including the full set of covariates in X made no big difference for the estimated treatment effects.

If we assume that wages are independently and identically distributed across individuals with conditional variance σ^2, it follows that (3) is heteroskedastic with conditional variance $\frac{\sigma^2}{n_j}$, where n_j is the number of employees in firm j. Estimation of (3) by means of ordinary least squares (OLS) would give unbiased and consistent estimators, but the OLS estimators are not efficient when the outcome variable is firm level average wages. The asymptotically efficient linear estimator is in this case the generalized least squares estimator, where (3) is weighted by the number of employees in firm j, known as weighted least squares (WLS). That is, the equation that is estimated by WLS is a transformed version of (3), where both sides of the equation has been multiplied with the square root of the number of employees in firm j. We use the number of employees in the base year (2005) to avoid the endogeneity problem that could be associated with the alternative of using the number of employees for each year t.

Turning to the firm level fixed effects specification, the equation to be estimated is equation (2), with $I_j \cdot d_t = D_{jt}$ and with industry specific year effects:

$$\bar{w}_{jkt} = \alpha_j + \lambda_{kt} t + \delta \left(I_j \cdot d_t \right) + X_{jt}^{'} \beta + \varepsilon_{jkt}, \tag{4}$$

where the dummies for geographical location are left out due to lack of variation across time periods. The weighting strategy described above is also applied to equation (4).

6.1 Initial results

We start by estimating (3) and (4) with the treatment dummy d_t taking the value $\frac{1}{2}$ for 2006 and 1 for the years 2007 through 2009. With the outcome variable $\bar{w}_{j(s)t}$ being defined as the natural log of average wages in firm j (in group s) at time t, $\hat{\delta}$ should be interpreted as *the estimated percentage difference in wage change (after vs before reform) between treated and non-treated firms.* Strictly negative treatment effects would thus indicate that there has been some degree of cost sharing between firms and employees, and to conclude that the complete costs are shifted onto employees we would require $\hat{\delta} < -0.026$, as the costs of a minimum requirement OP is approximately 2.6 percent of wages.

Table 3 shows results from OLS and WLS on (3) and WLS on (4), hereafter referred to as FE. The unweighted OP-group level fixed effects estimate of δ (Column 1) is a non-significant -0.3 per cent, the WLS estimate is a non-significant -0.5 per cent, whereas the firm level fixed effects estimate is precisely estimated at -1.3 per cent. This would indicate that firms were only able to shift about half (-1.3/-2.6) the costs onto their employees.

Table 3: Estimation results for equation (3) and (4)

	OLS	WLS	FE
Treated	-0.00316	-0.00507	-0.0131^{***}
	(0.00265)	(0.00329)	(0.00318)
DC firm	-0.0196^{***}	0.00673^{***}	
	(0.00152)	(0.00130)	
No OP firm	-0.0459^{***}	-0.0233^{***}	
	(0.00159)	(0.00193)	
Year×Industry	x	x	x
N	87210	87210	86652
R^2	0.593	0.677	0.650

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. OLS standard errors are heteroskedasticity robust. The dependent variable is log average earnings at the firm level, and the treatment indicator *Treated* takes the value .5 for No OP firm observations in 2006 and 1 for No OP firm observations from 2007 through 2009 and is zero otherwise. Additional controls are a cubic polynomial in the mean age of the employees, the fraction of male/immigrant/highly educated employees, and county dummies.

The results in Table 3 reveal two significant differences across specifications that deserve close attention. One is the fact that the firm level fixed effects average treatment effect is much larger in magnitude than the OLS and WLS treatment effects. This could reflect that low-wage firms are more likely to exit the sample than high-wage firms, and that the correlation between wage level and the exit probability is stronger for No OP firms than for the other two groups of firms. Such a relationship would be picked up by the firm level fixed effects, whereas OLS and WLS would tend to be biased towards zero, as the exits of low-wage firms would drive up the average wage level of the firms remaining in the sample. A probit model for the probability of being out of the sample in 2009 supports this explanation (see Table A1 in the Appendix). A second major difference is that the DC firm fixed effect changes from negative to positive when weights reflecting firm size in 2005 are added to equation (3). This could be indicative of very high wages or wage growth in some of the bigger comparison group firms than in other firms in the sample.

As neither of these differences are well accounted for by the available firm level characteristics when regressions are run on the full sample[8] it may well be that we are faced with a problem related to a lack of overlap in the covariate distributions of treatment and comparison group firms. One possible way around such a problem is to trade off some external validity against the benefits of working with a sample of more comparable firms. This path is explored in the following section.

6.2 Pre-screening based on the probability of treatment

When firms are very different in terms of observable characteristics there is a potential for selection bias even when these observables are included as controls in the regressions, and the assumption of common counter-factual trends for treatment and comparison group firms may be called into question. To cope with such a lack of overlap in the covariate distribution between treatment groups, we will follow Crump et al. (2009) in using the estimated probability of treatment as a basis for systematic pre-screening; They show that discarding all units with estimated propensity scores outside the range $[0.1, 0.9]$ may reduce the bias and decrease the variances of average treatment effect estimators. Estimated propensity scores from a probit model of the propensity to have an OP in the base year (2005) are depicted in Figure 3, for the full sample and for a restricted sample with $pscore \in [0.1, 0.9]$. Average marginal effects from the probit model

[8]Estimation results for (3) when the covariates in X_{jst} are included in a step-wise manner are shown in Table A2 in the Appendix. The OLS estimate of the treatment effect is fairly stable across specifications and never significantly different from zero. I have also estimated versions of (3) and (4) with four firm size specific treatment effects. Estimation results for these specifications gave no clear indications of heterogeneities with respect to firm size, and are thus not reported. Also unreported are results from specifications in which interactions with the treatment indicator and the right hand side variables were included, as these revealed no clear patterns, but rather underlined the discrepancies across specifications.

are given in Table A3 in the Appendix. The shape of the densities in Figure 3 suggest that there are characteristics with a very strong and positive relation with the propensity to have an OP, whereas there are no characteristics with predictive power of the same magnitude for the probability of not having an OP. This is confirmed by the estimates in Table A3, with firm size being the single most important predictor of having an OP.

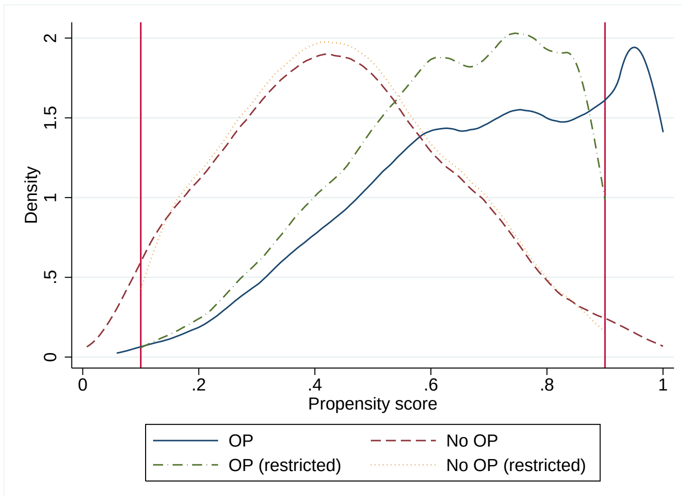

Figure 3: Estimated Kernel densities of propensity scores by OP status (2005). Vertical lines at 0.1 and 0.9. The restricted samples exclude firms with estimated propensity scores outside the range $[0.1, 0.9]$.

Table 4 shows estimated treatment effects from (3) and (4) for the full sample (Columns 1-3), for a sample of firms with estimated propensity scores below 0.9 (Columns 4-6), and for one consisting of firms with propensity scores within the range $[0.1, 0.9]$ (Columns 7-9).[9] First, we note that the OLS point estimate is reduced by about 22 per cent when the high propensity score firms are removed from the sample, and the standard error increases by about 5 per cent. WLS and FE point estimates are reduced by about 27 and 62 per cent, respectively, and standard errors are down by about 16 and 14 per cent. Secondly, the estimated OP-group fixed effect for DC-firms is now negative regardless of whether weights

[9] A small number of firms belonging to the Electricity, gas and water supply industry are excluded from all three samples in Table 4, as they were too few for the industry specific time trends to be properly identified. This exclusion had only minor impacts on the other estimated coefficients.

are used in the estimation. In sum, there are only modest differences across the three specifications when the high propensity score firms are excluded, and the fact that the average treatment effect estimator is now much less sensitive to the choice of specification makes a good case for proceeding with a reduced sample for the remainder of the analysis. Although very little is changed when the low propensity score firms are also excluded (Columns 7-9 vs Columns 4-6), we will follow Crump et al. (2009) and proceed with the sample consisting of firms with estimated propensity scores within the range $[0.1, 0.9]$.

Descriptive statistics for firms in the restricted sample for the years 2002 (pre-reform) and 2006 (the year of implementation) are given in Table 5. Problems related to a lack of overlap in the covariate distributions should be less likely for this sample than for the full sample, but might still be a concern. Hence I have reported the normalised difference for each of the covariates, defined as the difference in averages by treatment status, divided by the square root of the sum of the variances (Column 5 and 6). Imbens and Wooldridge (2009) suggest as a rule of thumb that linear regression methods might be sensitive to the functional form assumption if the normalised difference exceeds one quarter. This is the case for two of the variables: The number of employees and the mean age of the employees. The former enters only as weights (2005 values) in the regressions, and we saw in Table 4 how the inclusion of these weights has minor impacts on the estimated coefficients when regressions are run on the restricted sample. Mean age is included as a cubic polynomial in all regressions.

Although there are still differences in observed covariates between NoOP and OP firms, these do not appear to have changed dramatically over the observation period. Had there been substantial changes in the differences in the *observable* characteristics of treatment and comparison group firms over time, one might as well have been concerned about *unobserved* compositional changes. Such a concern would have called the difference-in-differences strategy into question.

Table 4: Estimation results for equation (3) and (4): Different samples

	Full sample			Sample with $pscore \leq 0.9$			Sample with $pscore \in [0.1, 0.9]$		
	OLS	WLS	FE	OLS	WLS	FE	OLS	WLS	FE
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
treated0609	-0.00337	-0.00523	-0.0132***	-0.00262	-0.00380	-0.00503*	-0.00296	-0.00404	-0.00514*
	(0.00265)	(0.00330)	(0.00317)	(0.00279)	(0.00276)	(0.00271)	(0.00280)	(0.00277)	(0.00272)
DC firm	-0.0194***	0.00679***		-0.0238***	-0.0130***		-0.0238***	-0.0130***	
	(0.00152)	(0.00130)		(0.00166)	(0.00147)		(0.00166)	(0.00148)	
No OP firm	-0.0457***	-0.0232***		-0.0481***	-0.0363***		-0.0480***	-0.0362***	
	(0.00159)	(0.00193)		(0.00172)	(0.00168)		(0.00172)	(0.00169)	
Year×Industry	x	x	x	x	x	x	x	x	x
N	86948	86948	86800	72380	72380	72235	71860	71860	71725
R^2	0.592	0.677	0.650	0.575	0.584	0.531	0.576	0.584	0.531

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. OLS standard errors are heteroskedasticity robust. The dependent variable is log average earnings at the firm level, and the treatment indicator *treated0609* takes the value .5 for 2006 and 1 for 2007-2009 for No OP firms and is zero otherwise. Additional controls are a cubic polynomial in the mean age of the employees, the fraction of male/immigrant/highly educated employees, and county dummies.

Table 5: Descriptive statistics for the restricted sample: $pscore \in [0.1, 0.9]$

	NoOP (means)		**Diff. in means** *NoOP - OP*		**Normalised diff.** *NoOP - OP*	
	2002	2006	2002	2006	2002	2006
Variable						
No. of employees	17.9578	18.8988	-8.9385	-7.1130	-0.2966	-0.3107
ln(Mean earnings)	12.6841	12.7746	-0.0600	-0.0704	-0.1958	-0.2057
Fraction of males	0.7849	0.7770	0.0371	0.0406	0.1103	0.1172
Avg. age of employees	39.2273	40.5319	-2.0179	-2.0651	-0.3106	-0.3128
Avg. tenure	5.9244	6.6206	-0.8684	-0.8656	-0.1839	-0.1751
Fraction highly educated	0.1431	0.1622	-0.0591	-0.0608	-0.1920	-0.1876
Fraction of immigrants	0.0550	0.0708	0.0036	0.0108	0.0273	0.0719
Industry[a]						
Manufacturing	0.2171	0.2077	-0.0421	-0.0439	-0.0700	-0.0739
Construction	0.2443	0.2348	0.1078	0.0995	0.1959	0.1827
Wholesale	0.2549	0.2525	-0.0173	-0.0119	-0.0278	-0.0192
Real estate	0.1296	0.1471	-0.0123	-0.0032	-0.0255	-0.0064
County[a]						
Akershus	0.0981	0.0928	-0.0012	-0.0069	-0.0029	-0.0166
Oslo	0.1659	0.1684	-0.0325	-0.0295	-0.0597	-0.0540
Rogaland	0.0863	0.0922	0.0008	0.0064	0.0020	0.0159
Hordaland	0.0942	0.0874	0.0084	-0.0045	0.0207	-0.0112
Sample size		*No OP*		*OP*		
		2002	2006	2002	2006	
		2,538	3,331	3,684	4,244	

[a] Same selection as in Table 2.

7 Refinements

7.1 Dynamic adjustments

As the case for mandatory occupational pensions was put forward already in 2002 and a settlement was reached in 2004, it may well be that the reform had effects on wages years before it formally entered into force January 1 2006. To allow for such anticipatory effects I have estimated versions of (3) and (4) with nine treatment dummies, one for each year starting from 2001; $D_{s\tau}$ now takes the value 1 for NoOP-firms in all years except 1999 and 2000. The difference-in-differences specification may now be written as

$$\bar{w}_{jkst} = \alpha + \gamma_s \mathit{OP\text{-}group}_s + \lambda_{kt} t + \sum_{\tau=2001}^{2009} \delta_\tau D_{s\tau} + X_{jt}^{'} \beta + \varepsilon_{jkst}, \qquad (5)$$

and the nine difference-in-differences estimators are defined as

$$\delta_\tau =$$

$$\begin{aligned}&\left\{E\left[\bar{w}_{jkst'}|s = NoOP, t' = \tau, X_{jt'}\right] - E\left[\bar{w}_{jkst}|s = NoOP, t \leq 2000, X_{jt}\right]\right\}\\ &- \left\{E\left[\bar{w}_{jkst'}|s = OP, t' = \tau, X_{jt'}\right] - E\left[\bar{w}_{jkst}|s = OP, t \leq 2000, X_{jt}\right]\right\},\end{aligned}$$

for $\tau = 2001, 2002, \ldots, 2009$. The fixed effects specification is modified accordingly.

Estimation results for equation (5) are given in Table 6. First, we note that the estimated treatment effects for the first three years are close to zero for all three specifications, meaning that the differences in wage levels between treatment and control group firms are now well accounted for by the covariates in X. Judging from the estimated treatment effects for 2004 and onwards it appears that the reform did affect wages prior to its formal implementation, although the effects were rather moderate, ranging from about -0.7 to about -1.2 per cent (FE). Based on these observations it appears that firms were holding back on wages already from the year in which a settlement was reached on mandatory occupational pensions (2004), and that the cost sharing peaked at nearly 50% (-1.2/-2.6) of the costs of a minimum requirement pension shifted onto workers in 2009.

I have also estimated equation (5) with a treatment effect also for the year 2000. The overall pattern is similar to the one seen in Table 6 and Figure 4, but the estimated treatment effects are smaller in magnitude and less precisely estimated when the tenth treatment effect dummy is included. Finally, when regressions are run on a balanced panel, that is, when only firms having at least ten full-time, full-year employees for all of the 11 years of observation are included, none of the estimated treatment effects are significantly different from zero. This might indicate that cost sharing has mainly taken place in not-so-stable firms.

Table 6: Estimation results for equation (5): Dynamic adjustments

	OLS	WLS	FE
treated01	-0.000187	0.00203	-0.000663
	(0.00533)	(0.00626)	(0.00309)
treated02	0.000691	0.00241	-0.00131
	(0.00508)	(0.00583)	(0.00310)
treated03	-0.00427	-0.00248	-0.00445
	(0.00497)	(0.00562)	(0.00328)
treated04	-0.0109**	-0.00918*	-0.00694**
	(0.00493)	(0.00557)	(0.00350)
treated05	-0.0171***	-0.0147***	-0.00990***
	(0.00481)	(0.00538)	(0.00353)
treated06	-0.0107**	-0.00943	-0.00943**
	(0.00508)	(0.00578)	(0.00393)
treated07	-0.00824	-0.00498	-0.00561
	(0.00523)	(0.00583)	(0.00402)
treated08	-0.00825	-0.00834	-0.0102**
	(0.00528)	(0.00591)	(0.00428)
treated09	-0.00948*	-0.0107*	-0.0123***
	(0.00525)	(0.00594)	(0.00452)
DC firm	-0.0239***	-0.0130***	
	(0.00166)	(0.00180)	
No OP firm	-0.0419***	-0.0318***	
	(0.00334)	(0.00381)	
Year×Industry	x	x	x
N	71860	71860	71725
R^2	0.576	0.584	0.531

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. OLS standard errors are heteroskedasticity robust. The dependent variable is log average earnings at the firm level, and the treatment indicators *treated01-treated09* take the value 1 for No OP firms in each of the years 2001-2009 and are zero otherwise. Additional controls are a cubic polynomial in the mean age of the employees, the fraction of male/immigrant/highly educated employees, and county dummies.

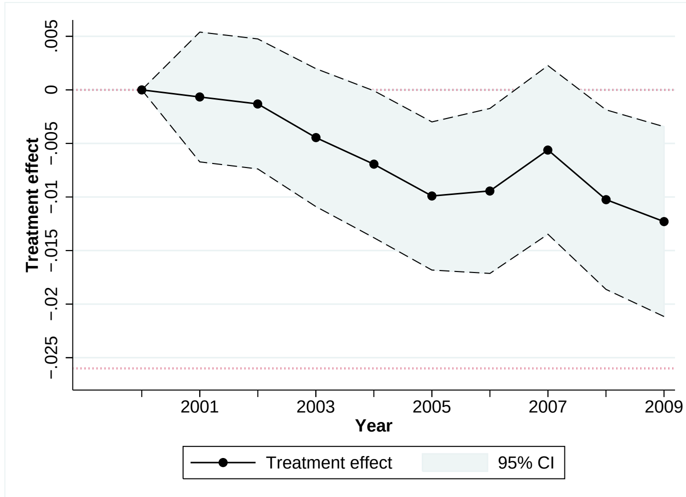

Figure 4: Estimated treatment effects from a model with firm level fixed effects that allows for pre- and post-reform effects, estimated on the sample with $pscore \in [0.1, 0.9]$. The dependent variable is log average earnings at the firm level.

7.2 The roles of labour unions and centralised negotiations

Knowing that the idea of introducing mandatory OPs in Norway was first advocated by labour unions, it seems reasonable to suspect that the degree of cost sharing may vary with the unions' influence on wages. According to Summers et al. (1993) and Ooghe et al. (2003) one should expect to see more cost sharing in economies where the wage setting is highly centralised, if unions are more likely than individuals to recognize the link between current contributions and future benefits. This line of reasoning is also supported by Alesina and Perotti (1997), who present a theoretical model that predicts a hump-shaped relationship between the degree of centralisation and the degree of shifting of labour taxation.

It is important to note, however, that both Summers et al. and Alesina and Perotti meant to explain cross-country differences and did not claim that the mechanisms at work within a country with a given level of centralisation are the same as those that are crucial for explaining differences between countries with different levels of centralisation. In central negotiations, the parties will neither know whether firms offer pensions, nor be able or willing to set wage increases that vary with current pension status.[10] Also, it might be that the impact of unions on the economic incidence of mandated OPs, for which there is a direct link between contributions and future benefits, is very different from their impact on the incidence of mandated social security contributions (Summers et al.) or general labour taxation (Alesina and Perotti).

In Norway there are central tariff negotiations between the biggest employer and employee organisations. A majority of the firms also take part in local negotiations, but some base their wage setting on central negotiations only, while others do not take part in neither central nor local negotiations.[11] If local union leaders are less concerned with economy wide considerations than the central leaders/unions, one would expect to observe a lesser extent of cost sharing where there is some local negotiations and where there are unions with high bargaining power (high share of unionised workers) than in activities where local unions are weak, either because their bargaining power is low or because local bargaining takes place at the individual level only.

Before moving on to the econometric analysis we will fix ideas by imagining an economy where both the influence of central negotiations and the bargaining power of local unions are characterised by binary variables; $Central$ and $UnionShare$. By combining the two we arrive at four sub-markets with different wage setting mechanisms, as described in Table 7. Markets with $Central = 1$ are those in which wages are influenced by central but not local negotiations, and markets with $UnionShare = 1$ are those in which all workers are unionised. When both variables are zero, that is, when there is a role for local unions to play, but these have no bargaining power, we would expect firm managers to take advantage of their relative bargaining power to make sure that costs are fully shifted onto the employees (panel A). In panel B, where local unions have high bargaining power, the firms will end up bearing the entire cost of

[10] An investigation of documents related to the central agreements before and after the introduction of mandatory OPs confirms that such discrimination did not take place.

[11] Cf. Løken and Stokke (2009) for a comprehensive overview of Norwegian labour relations.

the mandate (no shift), whereas the combination of local unions with no bargaining power and no local negotiations (panel C) would imply full shift. The outcome in the fourth sub-market (panel D) is indeterminate, as it is unclear how the local unions can exercise their bargaining power when there are no local negotiations.

Table 7: Implications of various wage setting mechanisms

	$UnionShare = 0$	$UnionShare = 1$
$Central = 0$	*A. Full shift*	*B. No shift*
$Central = 1$	*C. Full shift*	*D. Indeterminate*

To test these hypotheses we will use activity-specific union densities as a proxy for the unions' bargaining power in local negotiations: The variable $UnionShare_j$ is defined as the fraction of workers who deduct union membership fees from their pay check (which is the common way to pay membership fees among Norwegian union members) in firm j's three-digit NACE industry (activity). The use of union density as a proxy for the ability of unions to influence wages is supported by establishment-level evidence from Norway, documenting a strong relationship between union membership shares and individual wages (see Barth, Raaum, and Naylor (2000)). Another proxy for the unions' influence on the wage setting is the variable $Central_j$, which is constructed from a large firm level survey conducted in 2003 (Arbeids- og Bedriftsundersøkelsen (ABU 2003)). $Central_j$ is the fraction of firm managers in firm j's three-digit NACE industry who answered that wages are influenced by the outcomes of central but not local negotiations. $UnionShare$ is available for all 193 activities in the restricted sample, whereas $Central$ is available for 150 activities. Some descriptive statistics for each of the three groups of firms are given in Table 8, and a scatter plot of $Central$ against $UnionShare$ in Figure 5.[12]

[12] I am indebted to Bernt Bratsberg for providing me with the union membership series and

Table 8: Union membership and negotiations - summary statistics (2005)

	n	Mean	Std. Dev	Min	Max
UnionShare					
DB-firms	2,219	0.378	0.181	0.038	0.880
DC-firms	2,556	0.376	0.180	0.038	0.880
NoOP-firms	3,996	0.340	0.159	0.038	0.884
Central					
DB-firms	2,159	0.165	0.235	0	1
DC-firms	2,528	0.144	0.208	0	1
NoOP-firms	3,894	0.169	0.230	0	1
Corr $(UnionShare, Central) = -0.154$					

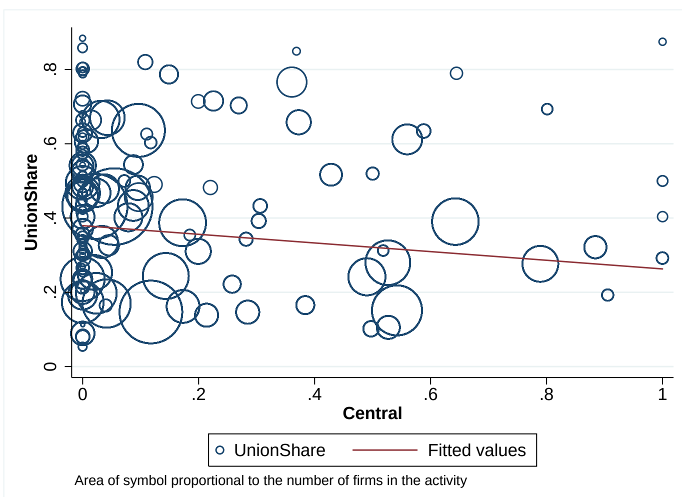

Figure 5: Scatter plot of *Central* against *UnionShare* for No OP firms.

The baseline specification is now modified in two ways: (i) To take account of the dynamic adjustments revealed in the preceding section the treatment indicator D_{st} now takes the value 0.5 for NoOP-firm observations from 2004 through 2006 (pre-reform), 1 for NoOP-firm observations from 2007 through 2009 (post-reform), and zero otherwise.[13] (ii) The treatment indicator is interacted with each of the variables $UnionShare_j$ and $Central_j$;

$$\begin{aligned}\bar{w}_{jkst} &= \alpha + \gamma_s OP\text{-}group_s + \lambda_{kt} t + \delta D_{st} \\ &\quad + \xi\left(UnionShare_j \cdot D_{st}\right) + \psi\left(Central_j \cdot D_{st}\right) + X_{jt}^{'}\beta + \varepsilon_{jkst}, \quad (6)\end{aligned}$$

where the respective means are subtracted from $UnionShare_j$ and $Central_j$. Identification of the total treatment effects now relies on variation across three-digit industries within two-digit industries. This may sound a bit far-fetched, but according to Bratsberg and Raaum (2010) there is substantial variation in industrial relations and wage setting institutions across activities in the construction industry, which is the industry to which 23% of the NoOP-firms in the sample belong.

Estimation results for equation (6) estimated on the reduced sample are given in Table 9, first with the *UnionShare* interaction only (Columns 1-3), then with the *Central* interaction only (Columns 4-6), and finally with both interactions included (Columns 7-9), and Table 10 gives total treatment effects at different values of the two variables based on the estimates in Column 8 (WLS). We first note that more appears to be shifted onto employees in activities with low shares of unionised workers, that is, where unions are likely to have relatively low bargaining power in local negotiations. It also seems to be the case that more is shifted in activities where wages being determined exclusively by the outcomes

the data on types of negotiations.

[13]The details of the Act were made public December 21 2005 and had economic effects for employees only from July 2006. These facts give reason to expect only partial treatment for the years 2004 – 2006, and are the rationale for letting the treatment indicator take the value 0.5 for these years. It turns out, however, that the results that follow are not sensitive to the precise definition of the treatment indicator: Testes with one single indicator taking the value 1 for the years 2004 – 2009 and with two indicators taking the value 1 for 2004 – 2006 and 2007 – 2009, respectively, gave results that were practically identical to the reported results.

of central negotiations is more common. There are only minor changes in the estimated treatment effects when both $UnionShare$ and $Central$ interactions are included (Columns 7-9 compared to 1-3 and 4-6).

To perform an informal test of the predictions of Table 7 I have computed WLS total treatment effects at the minimum and maximum values of $UnionShare$ and $Central$ in accordance with each of the four panels. This gives point estimates (standard deviations) of -0.026 (0.0065), 0.027 (0.0086), -0.047 (0.0088) and 0.005 (0.0131) for Panel A, B, C and D respectively, and does not reject our initial hypotheses. Although one should bear in mind that the predictive power of the linear regression model far away from the centre of the covariate distribution may be questionable, the results in this section seem to be in line with our prior expectations: More appears to be shifted onto employees in activities where wages are determined by the outcomes of central negotiations, whereas less is shifted onto employees in activities with high shares of unionised workers.

Table 9: The effects of union membership shares and central negotiations

	UnionShare			*Central*			*UnionShare* & *Central*		
	OLS	WLS	FE	OLS	WLS	FE	OLS	WLS	FE
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
DC firm	-0.0240***	-0.0133***		-0.0268***	-0.0162***		-0.0268***	-0.0162***	
	(0.00166)	(0.00148)		(0.00168)	(0.00149)		(0.00168)	(0.00182)	
No OP firm	-0.0459***	-0.0342***		-0.0454***	-0.0337***		-0.0456***	-0.0339***	
	(0.00195)	(0.00195)		(0.00199)	(0.00199)		(0.00199)	(0.00226)	
Treated	-0.00748**	-0.00793***	-0.00703**	-0.00856***	-0.00907***	-0.00814**	-0.00765**	-0.00813**	-0.00755**
	(0.00302)	(0.00303)	(0.00310)	(0.00305)	(0.00306)	(0.00317)	(0.00306)	(0.00344)	(0.00315)
UnionShare	-0.0268***	-0.0237***		-0.0353***	-0.0314***		-0.0436***	-0.0392***	
	(0.00661)	(0.00554)		(0.00654)	(0.00560)		(0.00694)	(0.00769)	
Treated×UnionShare	0.0531***	0.0604***	0.0489***				0.0494***	0.0577***	0.0514***
	(0.0123)	(0.0118)	(0.0169)				(0.0129)	(0.0143)	(0.0183)
Central				-0.116***	-0.101***		-0.117***	-0.102***	
				(0.00392)	(0.00369)		(0.00394)	(0.00444)	
Treated×Central				-0.0246**	-0.0286***	-0.00818	-0.0193**	-0.0217**	-0.00227
				(0.00797)	(0.00822)	(0.0103)	(0.00821)	(0.00934)	(0.0108)
N	71860	71860	71725	69175	69175	69041	69175	69175	69041
R^2	0.576	0.584	0.531	0.588	0.596	0.530	0.589	0.596	0.530

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. OLS standard errors are heteroskedasticity robust. The dependent variable is log average earnings at the firm level, and the treatment indicator *Treated* takes the value .5 for No OP firm observations from 2004 through 2006, 1 for No OP firm observations from 2007 through 2009, and zero otherwise. Additional controls are a cubic polynomial in the mean age of the employees, the fraction of male/immigrant/highly educated employees, and county dummies.

Table 10: Total treatment effects (based on Table 9 (Column 8))

	UnionShare		**Central**	
p10	-0.0207472***	(0.004448)	-0.0046892	(0.0037284)
p20	-0.0194098***	(0.0042448)	-0.0046892	(0.0037284)
p25	-0.0180878***	(0.0040605)	-0.0051436	(0.0036576)
p30	-0.0168169***	(0.0039011)	-0.0051477	(0.003657)
p40	-0.0138573***	(0.0036095)	-0.0055095	(0.0036072)
p50	-0.0087303**	(0.0034311)	-0.0058757*	(0.003563)
p60	-0.0043458	(0.0036402)	-0.0067774*	(0.0034822)
p80	-0.0028095	(0.0037849)	-0.0130343***	(0.0040554)
p90	0.0020448	(0.0044296)	-0.0164698***	(0.0049967)
p95	0.0074976	(0.005386)	-0.0186574***	(0.0057152)
p99	0.0120965*	(0.0063053)	-0.0238994***	(0.0076353)

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The percentiles in the first column are those of the distribution of each of the variables *UnionShare* and *Central* for NoOP-firms.

8 Conclusion

The purpose of this paper has been to provide new insights about the economic incidence of mandatory occupational pensions, and more generally about the offset between pensions and wages. To overcome simultaneity biases caused by the joint determination of pensions and wages, the recent introduction of mandatory OPs in Norway has been exploited as a source of exogenous variation in pension coverage. Various difference-in-differences models have been estimated on a large panel of Norwegian private sector firms covering the years 1999 through 2009, and estimation results indicate that firms have only been able to shift parts of the costs onto employees: The largest estimated average treatment effect from the most flexible fixed effects model indicates that less than 50% of the costs of a minimum requirement OP was shifted onto workers in 2009, three years after the reform was formally implemented. The results also reveal considerable heterogeneities across activities with different wage setting mechanisms: More appears to be shifted onto employees in activities where wages are influenced by central, but not local negotiations, and in activities with low shares of unionised workers.

Although the documented heterogeneities with respect to the influence of local unions and central negotiations on the wage setting may be given a logically sound interpretation, they are not necessarily in line with predictions from standard models of bargaining. These findings may therefore deserve a more thorough theoretical analysis. Also, given that the average treatment effects suggest that a significant proportion of the costs are borne by employers, one should expect to see responses also in employment, prices or profits. A complete analysis of employment effects should take account of possible compositional changes and changes in employment levels, which would require an analytical framework different from the one applied in this paper. Sufficiently detailed data on prices and profits are not readily available, and hence an assessment of price and profit effects of the mandate must also be left for further research.

References

Alesina, A. and R. Perotti (1997). The welfare state and competitiveness. *The American Economic Review 87*(5), pp. 921–939.

Angrist, J. D. and J.-S. Pischke (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.

Azémar, C. and R. Desbordes (2010). Who ultimately bears the burden of greater non-wage labour costs? Discussion Paper 2010-02, Department of Economics, University of Glasgow.

Barth, E., O. Raaum, and R. Naylor (2000). Union wage effects: Does membership matter? *Manchester School 68*(3), 259–75.

Bratsberg, B. and O. Raaum (2010). Immigration and wages: Evidence from construction. CReAM Discussion Paper Series 1006, Centre for Research and Analysis of Migration (CReAM), Department of Economics, University College London.

Crump, R. K., V. J. Hotz, G. W. Imbens, and O. A. Mitnik (2009). Dealing with limited overlap in estimation of average treatment effects. *Biometrika 96*(1), 187–199.

Gruber, J. (1994). The incidence of mandated maternity benefits. *The American Economic Review 84*(3), 622–641.

Gustman, A. L., O. S. Mitchell, and T. L. Steinmeier (1994). The role of pensions in the labor market: A survey of the literature. *Industrial and Labor Relations Review 47*(3), 417–438.

Hernæs, E., J. Piggott, O. L. Vestad, and T. Zhang (2011). Labour mobility, pension portability and the lack of lock-in effects. *SSRN eLibrary*.

Hernæs, E., J. Piggott, T. Zhang, and S. Strøm (2010). Occupational pensions, tenure, and taxes. *Journal of Pension Economics and Finance FirstView*, 1–22.

Imbens, G. W. and J. M. Wooldridge (2009). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature 47*, 5–86(82).

Iskhakov, F. (2008). Pension reform in Norway: Evidence from a structural dynamic model. Memorandum 14/2008, Department of Economics, University of Oslo.

Lazear, E. P. (1981). Agency, earnings profiles, productivity, and hours restrictions. *The American Economic Review 71*(4), 606–620.

Løken, E. and T. A. Stokke (2009). Labour relations in Norway. Fafo-report 2009:33.

Midtsundstad, T. and J. M. Hippe (2005). Pensjoner i arbeidsmarkedet. In H. Torp (Ed.), *Nytt arbeidsliv. Medvirkning, inkludering og belønning*, pp. 229–249. Gyldendal Akademisk.

Montgomery, E., K. Shaw, and M. E. Benedict (1992). Pensions and wages: An hedonic price theory approach. *International Economic Review 33*(1), pp. 111–128.

Ooghe, E., E. Schokkaert, and J. Flechet (2003). The incidence of social security contributions: An empirical analysis. *Empirica 30*(2), 81–106.

Schiller, B. R. and R. D. Weiss (1980). Pensions and wages: A test for equalizing differences. *The Review of Economics and Statistics 62*(4), pp. 529–538.

Smith, R. S. (1981). Compensating differentials for pensions and underfunding in the public sector. *The Review of Economics and Statistics 63*(3), pp. 463–468.

Summers, L., J. Gruber, and R. Vergara (1993). Taxation and the structure of labor markets: The case of corporatism. *The Quarterly Journal of Economics 108*(2), 385–411.

Summers, L. H. (1989). Some simple economics of mandated benefits. *The American Economic Review 79*(2), 177–183.

Veland, G. (2008). Tjenestepensjonsordningene i Norge. En undersøkelse av status og utviklingstrekk i privat sektor. Fafo-rapport 2008:23. Oslo: Fafo.

A Appendix

A.1 A probit model of $Pr($***Out of sample in 2009***$)$

Table A1: Average marginal effects from probit model

	Coefficient		Standard error
Treated	0.603***		(0.184)
ln(Mean earnings)	-0.0957***		(0.0275)
Treated×ln(Mean earnings)	-0.147***		(0.0358)
N	10385		
Dependent mean	0.258		
pseudo R^2	0.057	ll	-5593.1

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable = 1 if the firm is out of the sample in 2009. *Treated* takes the value 1 for No OP firms and is zero otherwise. Additional controls are a cubic polynomial in the mean age of the employees, the fraction of male/immigrant/highly educated employees, and dummies for industry affiliation and county.

Table A2: Estimation results for equation (3), stepwise inclusion of covariates

	OLS_I	OLS_II	OLS_III	OLS_IV	OLS_V	OLS_VI	OLS_VII	OLS_VIII
treated0609	-0.00275	-0.00456	-0.00370	-0.00248	-0.00250	-0.00472	-0.00434	-0.00316
	(0.00372)	(0.00335)	(0.00343)	(0.00322)	(0.00309)	(0.00307)	(0.00267)	(0.00265)
Year	x	x						
OP-group	x	x	x	x	x	x	x	x
Industry		x						
Year×industry			x	x	x	x	x	x
County				x	x	x	x	x
Fracmale					x	x	x	x
Meage						x	x	x
Frachi							x	x
Fracimm								x
_cons	12.72***	12.66***	12.66***	12.80***	12.60***	9.671***	10.08***	10.12***
	(0.00319)	(0.00304)	(0.00467)	(0.00461)	(0.00521)	(0.182)	(0.167)	(0.165)
N	87210	87210	87210	87210	87210	87210	87210	87210
R^2	0.120	0.295	0.297	0.393	0.432	0.438	0.588	0.593

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. OLS standard errors are heteroskedasticity robust.

A.2 A probit model of $Pr(OP = 1)$

Table A3: Average marginal effects from probit model

	Coefficient		Standard error
_Inace_1	-0.0986*		(0.0405)
_Inace_2	0.0940*		(0.0416)
_Inace_5	-0.0600***		(0.0146)
_Inace_6	-0.0205		(0.0125)
_Inace_7	0.0642*		(0.0258)
_Inace_8	-0.0566**		(0.0196)
_Inace_9	0.104**		(0.0330)
_Inace_10	-0.118***		(0.0172)
_Inace_11	-0.0760		(0.0429)
_Inace_12	0.0205		(0.0263)
_Inace_13	-0.0525		(0.0282)
lnnemp	0.241***		(0.00598)
_Icounty_1	-0.0444*		(0.0219)
_Icounty_2	-0.0194		(0.0166)
_Icounty_4	0.0330		(0.0246)
_Icounty_5	0.0261		(0.0241)
_Icounty_6	-0.0435*		(0.0209)
_Icounty_7	-0.0461*		(0.0215)
_Icounty_8	0.0238		(0.0260)
_Icounty_9	-0.0940**		(0.0355)
_Icounty_10	-0.0000714		(0.0266)
_Icounty_11	0.0128		(0.0174)
_Icounty_12	0.00876		(0.0173)
_Icounty_13	0.0134		(0.0296)
_Icounty_14	0.0247		(0.0207)
_Icounty_15	0.00275		(0.0203)
_Icounty_16	0.0944***		(0.0284)
_Icounty_17	0.0377		(0.0230)
_Icounty_18	0.0111		(0.0280)
_Icounty_19	-0.0644		(0.0427)
frachi	0.296***		(0.0237)
fracmale	-0.0578**		(0.0206)
fracimm	-0.217***		(0.0441)
meage	0.134		(0.0834)
meage2	-0.00246		(0.00202)
meage3	0.0000173		(0.0000162)
N	10365		
pseudo R^2	0.210	ll	-5513.7

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Dependent variable = 1 if the firm had an OP in 2005 (base year), 0 otherwise.