

Government
Institute for
Economic Research

Publications 60

Essays on labour demand and wage
formation

VATT PUBLICATIONS

60

Essays on labour demand and wage formation

Ossi Korkeamäki

ISBN 978-952-274-011-3 (nid.)
ISBN 978-952-274-012-0 (PDF)

ISSN 0788-4990 (nid.)
ISSN 1795-3332 (PDF)

Valtion taloudellinen tutkimuskeskus

Government Institute for Economic Research

Arkadiankatu 7, 00100 Helsinki, Finland

Email: etunimi.sukunimi@vatt.fi

Edita Prima Oy

Helsinki, January 2012

Graphic design: Niilas Nordenswan

Essays on labour demand and wage formation

Ossi Korkeamäki

Academic dissertation to be presented, by the permission of the Faculty of Social Sciences of the University of Helsinki, for public examination at Economicum, Lecture Room, Arkadiankatu 7, on February 3, 2012, at 12 noon.

Essays on labour demand and wage formation

Government Institute for Economic Research
VATT Publications 60/2012

Ossi Korkeamäki

Abstract

The first essay in this thesis is on gender wage differentials among manufacturing sector white-collar workers. The wage differential is decomposed into firm, job (within-firm) and individual-level components. Job-level gender segregation explains over half of the gap, while firm-level segregation is not important. After controlling for firm, job and individual characteristics, the remaining unexplained wage gap to the advantage of men is six per cent of men's mean wage.

In the second essay, I study how the business cycle and gender affect the distribution of the earnings losses of displaced workers. The negative effect of displacement is large, persistent and strongest in the lowest earnings deciles. The effect is larger in a recession than in a recovery period, and in all periods women's earnings drop more than men's earnings.

The third essay shows that the transition from steady employment to disability pension depends on the stringency of medical screening and the degree of experience-rating of pension costs applied to the employer. The fact that firms have to bear part of the cost of employees' disability pension costs lowers both the incidence of long sick leave periods and the probability that sick leave ends in a disability pension.

The fourth and fifth essays are studies on the employment, wage and profit effects of a regional payroll tax cut experiment conducted in northern and eastern Finland. The results show no statistically significant effect on any of the response variables.

Key words: Gender wage gap, gender segregation, displacement, earnings losses, disability pension, experience rating, payroll tax, tax incidence, labour demand

Tiivistelmä

Ensimmäisessä esseessä miesten ja naisten välinen keskipalkkojen ero jaetaan sukupuolisegregaatiosta johtuvaan osaan ja henkilökohtaisista ominaisuuksista johtuvaan osaan. Tulosten mukaan yksityisen sektorin toimihenkilöillä yritystason segregatio ei juuri ole osallisena palkkaeron syntyyn. Hiukan yli puolet kokonaiserosta selittyy naisten keskittymisellä matalapalkkaisiin töihin yritysten sisällä. Kun otetaan huomioon erot koulutuksessa, työkokemuksessa ja työn vaativuudessa on selittymätön palkkaero kuusi prosenttia miesten keskipalkasta.

Toisessa esseessä näytetään, että irtisanomisen vaikutus ansiotulojakaumaan on pitkäkestoinen ja voimakkain jakauman alimmissa desiileissä. Irtisanomisen vaikutus tulojakaumaan on paljon suurempi laman aikana kuin kasvuperiodilla ja negatiivinen vaikutus naisten tuloihin on kaikissa oloissa voimakkaampi kuin miehillä.

Kolmannen esseen tulokset osoittavat, että työntekijän terveydentilan lisäksi myös lääketieteellisten kriteerien tiukkuus ja yritysten omavastuu eläkkeen kustannuksista vaikuttavat työkyvyttömyyseläkkeelle jäämisen todennäköisyyteen. Vaikutus on kaksiosainen: omavastuu vähentää pitkiä sairauslomia ja kasvattaa työntekijöiden todennäköisyyttä palata takaisin töihin sairastumisen jälkeen

Neljännessä ja viidennessä esseessä tutkitaan Pohjois-Lapin ja Kainuun alueella toteutetun yritysten sotumaksuvapautuksen vaikutusta yritysten työllisyyteen, palkkasummaan ja voittoihin sekä työntekijöiden palkkoihin. Kokeilulla ei havaittu olevan tilastollisesti merkitsevää vaikutusta yhteenkään vastemuuttajaan.

Asiasanat: Sukupuolten palkkaero, segregatio, irtisanominen, työkyvyttömyyseläke, yritysten omavastuu, palkan sivukulut, työn kysyntä, verotuksen kohtaanto

Acknowledgements

The foreword or acknowledgements section of a thesis usually describes the process of writing the thesis: where it all began, who and what were the main influences along the way – sometimes even what it all was like. In my case, were I compelled to condense everything to the minimum, I could do it in four letters: VATT. The Government Institute for Economic Research provided the necessary resources and supervision. It tolerated delays when I stalled and then pushed me forward again. It might seem impersonal to give a first round of thanks to an institute, but I believe that VATT has been more than the sum of the persons that work there. For this I wish to thank the Directors General Reino Hjerppe, Seija Ilmakunnas and Aki Kangasharju.

My first supervisor and superior at VATT was Research Director Pasi Holm. He was followed by Seija Ilmakunnas, Heikki Räisänen, Roope Uusitalo, and for the last leg by Anni Huhtala. I thank them for their encouragement and for their unwavering belief in my ability to finish this project. If there was any doubt, they kept it to themselves and kept me going... and going. I might well be the last economist to have started his thesis in the previous millennium and actually complete it. Roope is also co-author of one of the articles and became my academic supervisor towards the end of this venture. I am grateful for his help and advice. I would also like to thank Kari Hämäläinen. In his own words I was “unsupervisable”, but his door was always open and quite often I dropped in for advice anyway.

If one only scans the contents of my thesis it is obvious to whom the most gratitude is due. Tomi Kyyrä is co-author of three of the papers. He is a great colleague and without him I doubt this thesis would exist. Jukka Appelqvist worked with us on the displacement paper and Antti Luukkonen took part in the gender wage differentials project. Thank you Jukka and Antti. There are many other colleagues I want to thank for their good company and actual help: Juha Tuomala, Matti Sarvimäki, Teemu Lyytikäinen and Tuomas Kosonen. Whenever I needed help with computers, I could rely on Raimo Hintikka’s swift response. Sari Virtanen was essential in getting this document printed and Andrew Lightfoot helped me with the language.

The other vitally important institute for my research was Statistics Finland. I wish to thank the persons there working for the research laboratory: Mika Maliranta, Satu Nurmi, Jouko Verho, Antti Katainen and Marjo Pyy-Martikainen.

I would like to thank my pre-examiners, Professors Rudolf Winter-Ebmer and Oskar Nordström Skans, for their comments.

The Academy of Finland, the Yrjö Jahnsso Foundation and Labour's Foundation provided financing for my studies and for completing this thesis.

On a more personal note, I would like to thank Janice Redman, Tarja Nyberg and the Sparrows. Thank you Hannele, Joonas and Matias.

Contents

1. Introduction	1
1.1 A gender wage gap decomposition for matched employer-employee data	1
1.2 A distributional analysis of earnings losses of displaced workers in an economic depression and recovery	3
1.3 Institutional rules, labour demand and retirement through disability programme participation	5
1.4 Employment and wage effects of payroll tax cut – evidence from a regional experiment	7
1.5 The Finnish payroll tax cut experiment revisited, or where did the money go?	9
2. A Gender wage gap decomposition for matched employer-employee data	13
2.1 Introduction	13
2.2 Methodological framework	16
2.2.1 The wage model	16
2.2.2 Decomposing the gender gap in pay	17
2.2.3 The fixed effects approach	18
2.2.4 The correlated random effects approach	19
2.2.5 Discussion	21
2.3 Data and descriptive statistics	23
2.3.1 The TT data and job classification	23
2.3.2 Sample statistics	25
2.4 Results	30
2.4.1 Wage regressions	30
2.4.2 Wage gap decompositions	33
2.4.3 Robustness of the results	36
2.4.4 Comparisons with findings from other studies	40
2.5 Conclusion	42
3. A distributional analysis of earnings losses of displaced workers in an economic depression and recovery	45
3.1 Introduction	45
3.2 Evaluation issues	47

3.2.1	Expected earnings losses	48
3.2.2	Distributional analysis	50
3.3	Data and sample construction	52
3.4	Descriptive evidence	55
3.4.1	Macroeconomic environment	55
3.4.2	Background characteristics	57
3.4.3	Empirical earnings distributions	59
3.5	Quantile displacement effects	62
3.5.1	Pre-displacement effects	71
3.5.2	Post-displacement effects	73
3.5.3	Robustness of the results	75
3.5.4	Comparisons to results from mean regressions	77
3.6	Concluding remarks	80
4.	Institutional rules, labour demand and retirement through disability programme participation	85
4.1	Introduction	85
4.2	Related literature	87
4.3	Institutional framework of Finland	91
4.3.1	Sickness and disability benefits	91
4.3.2	Experience rating of disability pension benefits	93
4.4	Data and descriptive evidence	95
4.4.1	Incidence of disability retirement	96
4.4.2	Outcome variables for analysis of transitions	97
4.4.3	Sample design for modelling transition rates	100
4.4.4	Raw transition rates	100
4.5	Determinants of transition rates	104
4.5.1	Individual characteristics	104
4.5.2	Strictness of medical criteria	107
4.5.3	Experience rating	108
4.5.4	Employment growth and excess turnover	112
4.6	Concluding remarks	115
5.	Employment and wage effects of a payroll tax cut – evidence from a regional experiment	119
5.1	Introduction	119
5.2	The experiment	121

5.3 Tax incidence and the Finnish wage bargaining system	124
5.4 Empirical strategy	125
5.5 Data	129
5.6 Results	131
5.6.1 Covariate balancing	131
5.6.2 Employment and wage sum responses to the regional payroll tax experiment	134
5.6.3 The effects by firm type	136
5.6.4 The effect on wages	137
5.7 Concluding comments	139
6. The Finnish payroll tax cut experiment revisited, or where did the money go?	143
6.1 Introduction and background	143
6.2 The experiment, target and comparison regions and firms	145
6.2.1 Finnish payroll taxes	146
6.2.2 Target and comparison regions used in the evaluation	147
6.2.3 Target and comparison firms	151
6.3 Data sets	154
6.4 Identification	155
6.5 Results	157
6.5.1 Where did the money go?	159
6.5.2 Effect on employment, wage sum and profits	161
6.5.3 Effect on wages and hours, individual wage records	163
6.5.4 Robustness checks	166
6.6 Discussion and some conclusions	171

1. Introduction

This thesis contains five microeconomic studies where I explore wage formation and labour demand in the Finnish labour market. All the essays are based on linked employer-employee data.

The first paper is an investigation into gender wage differentials. The focus is on to what extent these differentials arise from the segregation of the labour market into men's jobs and women's jobs on the one hand, and from segregation of men and women into different firms, and hence industries, on the other. The second essay is a distributional analysis of the earnings losses of displaced workers. The gender aspect is also present in the second study as clear differences are found in how the earnings of men and women respond to losing one's job. The second essay also shows that macroeconomic labour demand conditions have a huge impact on the individual's labour market success after displacement. The third study concerns the effects of experience-rating firms' disability pension costs on the incidence of disability pensions. There I show that the design of institutions involved with disability pensions has a strong effect on firms' labour demand decisions. In the fourth paper, I evaluate the employment and wage effects of a regional payroll tax cut. The last study widens the scope of the evaluation to include the effects of the tax cut on the profitability of firms.

This introduction continues with a short description of each of the studies. I discuss the main results and conclusions and try to highlight the contributions to our knowledge of the Finnish labour market. The papers are independent in their surveys of previous literature, and the methods and datasets are also thoroughly described. Hence, rather than constructing the theoretical background, building the apparatus for analysis and describing the relevant institutions, I aim at brevity in this introductory chapter.

1.1 A gender wage gap decomposition for matched employer-employee data¹

In this paper the gender wage gap is decomposed based on a correlated random effects model. The decomposition makes it possible to assess the extent to which the overall gap is attributable to gender segregation at the firm level, gender segregation into jobs within firms and within-job wage differentials. The data set comes from the Confederation of Finnish Industries and covers large and medium-sized manufacturing sector firms.

¹ Korkeamäki, O. – Kyyrä, T. (2006): A gender wage gap decomposition for matched employer-employee data. *Labour Economics*, vol. 13(5), pp. 611–638.

By explicitly modelling the firm and job effects, the approach also proves to be informative regarding the sources of lower pay in predominantly female firms and jobs. The data contains a detailed measure of job complexity that makes it possible to compare jobs within firms. In this study I present the results for manufacturing sector white-collar workers. The working paper version also contains the results for blue-collar workers.

The difference in mean wages between men and women is approximately 22 per cent of men's mean wage, to the advantage of men. Firm-level segregation accounts for only a small part, three percentage points, of the differential. At the firm level, no firm or firm-averaged worker characteristics are strongly associated with the gap. The difference in the share of women of a firm's employees between men and women explains only half of a percentage point of the total difference in wages.

The majority of the gender wage differential, a little over half, is attributed to the disproportionate concentration of women in lower-paying jobs within firms. High-paid managerial jobs are mainly occupied by men, which explains one fifth of the total gap. Among other types of jobs, men are concentrated in positions with higher skill requirements – that explains a tenth of the gap. Apart from differences in skills and managerial ability, these parts of the gap may reflect discrimination through differential access to higher-paying jobs, or they may result from gender differences in preferences. Although the reasons for the preponderance of women in lower-paying jobs remain a puzzle, our findings highlight the importance of equal opportunities in education, hiring, and promotion.

However, the results also suggest that predominantly female jobs pay lower wages than predominantly male jobs even if they are associated with a similar level of average education, tenure and job complexity. In other words, jobs of equal worth are differently rewarded depending on whether they are occupied by men or women. Of course, one can always speculate how accurate the job complexity variable is, but if we assume that this measure is reasonably good, the results would imply that policies like comparable worth should be considered. One third of the wage differential arising from the level of jobs within firms is due to the “femaleness” of the job and cannot be explained by other variables.

Finally, I found that within jobs women are paid some six per cent less than their equally qualified male co-workers. The unexplained within-job gap is higher among more educated and more experienced workers. Eliminating the sources of unexplained within-job wage differentials can directly account for a quarter of the overall gender gap in pay. On the one hand, this six per cent gap is a lot smaller than the often-quoted figure of 20 per cent drawn from gender averages; on the other hand, it is far from being insignificant. It is surprising that a gap of this size in hourly pay exists for equally qualified workers in the same job and

firm. The findings are in line with results from Sweden and Norway, but in Denmark the unexplained within-job wage gap is much larger – 14%.

The earliest gender pay differential studies concentrated mainly on personal characteristics, largely omitting the role of labour market segregation, whereas the first segregation-oriented papers used the female share at different levels of aggregation (industry, occupation, firm and job) as explanatory variables. In this study, the nested structure of the data is taken into account in the estimation and the results show that not doing so will give misleading results.

1.2 A distributional analysis of earnings losses of displaced workers in an economic depression and recovery²

The gender wage differential study is, strictly speaking, a descriptive cross-section data story. The analysis is done meticulously, even obsessively so, but the dynamics leading to a certain wage distribution could only be guessed at. In the second study of this thesis, I look at a specific event in the labour market, displacement from a firm, and try to track its effects on the distribution of earnings over a period of time.

I analysed the earnings losses owing to involuntary job loss among Finnish workers who became displaced during a period of depression (1992) or recovery (1997). These groups of displaced workers and the associated comparison groups were followed over an 11-year period beginning three years before and ending seven years after the year of possible displacement. A few years ago, the early 1990s Finnish recession could have been considered an extreme case and not likely to have much relevance in more normal times. Now that the financial crisis has dealt a serious blow to a number of economies, the results could give some guidance to its possible labour market consequences.

Using the quantile regression method, I estimated the effect of displacement at each decile of the earnings distribution. The findings from both periods suggest that 1) displaced workers suffer from substantial and persistent earnings losses, 2) women are subject to larger earnings losses than men, and 3) the effect of displacement is very heterogeneous, being much larger in the lower quantiles and implying a sharp increase in the earnings dispersion following displacement. The fourth finding comes from the comparison of the effects of displacement at different phases of the business cycle: in a recession the earnings distribution of

² Unpublished manuscript. Earlier working paper version: Korkeamäki, O. – Kyyrä, T. (2008): A Distributional analysis of displacement costs in an economic depression and recovery. VATT Discussion Papers 465, Government Institute for Economic Research.

the displaced workers falls far below the counterfactual distribution; in a recovery period only the lower half of the earnings distribution is affected.

The first finding is in accordance with results from the US labour market. The results from other European labour markets are more mixed in this respect. The second result is interesting given that most US studies have not found notable differences between women and men, whereas the gender aspect is given very little attention in European studies. One possible explanation for the larger earnings losses of women could be that in Finland, as in other countries, women are more frequently out of work for family reasons. That may induce employers to favour male employees when investing in managerial and professional skills. If so, and if such skills are generally transferable, i.e. not lost in job displacement, one can expect to find smaller earnings losses for displaced men. Taken together with the fact that managerial, professional and technical jobs are disproportionately held by men, employers' investment behaviour may lead to larger earnings losses for women. It remains unclear why earnings losses differ between men and women in Finland, but not in the US. It should be stressed that women's labour market position is quite different in Finland. On the one hand, the relatively generous maternity and parental leave schemes encourage career breaks, but on the other hand, public day care and school meals help the mothers of young children to work full-time if they want to. Moreover, part-time work among women in Finland is not very common and the labour force participation rate of women is relatively high.

The third finding, the heterogeneity in the displacement effect, is the main contribution of the paper and has important implications. First, the positive effect on earnings dispersion means that job loss does not only cause a significant decline in expected earnings, but also creates uncertainty about the level of future earnings. This suggests an additional welfare loss for risk-averse workers. This effect has typically been ignored in the discussion of displacement costs. Secondly, the large effect at the lower end of the distribution is consistent with the hypothesis that the relative importance of transferable individual-specific skills, which are not lost in job displacement, is larger for high-ability workers, who tend to populate the upper part of the conditional earnings distribution. Finally, the disproportionately large effect on the first two deciles implies that the effect on the expected earnings loss is in large part driven by an increased risk of joblessness and low-paid employment following job displacement. This means that job training and job replacement programmes targeted at unemployed job seekers, if effective in enhancing re-employment, can provide a means to reduce the average displacement cost.

By comparing the results from the two periods (displacement in 1992 or 1997), we found much larger earnings losses for those who lost their jobs during the depression period. Men (women) who were displaced in the middle of the depression had approximately 58% (65%) lower median earnings one year after

job loss and 15% (20%) lower median earnings seven years after job loss. Because of the exceptionally difficult labour market conditions, their earnings distribution as a whole remained below the counterfactual level through the end of the follow-up period. By contrast, job loss in the recovery period had a long-lasting effect only on the lower half of the distribution. For workers displaced in 1997 the effect on the median of the earnings distribution of men was 3.9% (women 7.1%) one year after being displaced. Seven years later the effect was still approximately 8% for women, whereas it did not differ from zero for men. These long-term losses do not vanish even if we account for income transfers.

1.3 Institutional rules, labour demand and retirement through disability programme participation³

The disability benefit scheme is one of the largest social security programmes in many countries, and therefore is of particular interest. In Finland, disability is the most common reason for early retirement, and disability expenditure accounted for some 3.5% of GDP in 2003, which was the third highest share in the EU after Sweden and Denmark. Disability enrolment rates of older employees vary strikingly across the European countries and the US. These cross-country differences cannot be explained by demographic or health-related factors. Over the past two or three decades, many countries have also experienced an expansion of disability benefit enrolment even though their ageing populations have become healthier. This is a serious concern given the common goal of inducing people to retire later. The widespread use of disability benefits as an early retirement instrument has been argued to be a particularly serious problem in Finland.

When job cuts are necessary, firms often offload their oldest employees first. If the health requirements for disability benefit eligibility are weak, early retirement via the disability scheme can be a useful strategy in effective downsizing, providing a way to reduce the workforce in a “soft” way. Some firms may also target dismissals at those employees with a high risk of disability. In doing so, the employer can avoid disability costs arising from the experience-rated contributions of disability pension benefits. Encouraging disability retirement could also be an attractive strategy for an employer wanting to change the composition of the workforce at a time of stable or growing employment when dismissals are difficult to justify.

Much previous empirical literature has been based on a simple labour supply framework in which an employee chooses whether to apply for disability

³ Korkeamäki, O. – Kyyrä, T. (2011): Institutional rules, labour demand and retirement through disability programme participation. *Journal of Population Economics*, forthcoming, available online.

benefits, while the employer has no role at all. Surprisingly little effort has been made to study the labour demand side. This essay aims to shed light on the relationships between labour demand, institutional factors and early retirement through disability programmes. I consider the importance of the labour demand side by examining the relationship between the growth and restructuring rates of an establishment and entries into disability. In addition, I assess the effectiveness of two policy instruments: the strictness of medical requirements for disability pension eligibility and the experience-rating of disability expenditure. The first determines the ease of access to disability pension benefits, whereas the latter places part of the costs of early retirement on the employer.

Transitions out of work to sick leave and disability retirement are modelled using matched employer-employee data for the Finnish private sector covering the years 1991–2005. The data set includes all active firms and employees can be tracked across all labour market states. To identify the role of institutional factors I exploit a law change that made the medical requirements for disability pension eligibility tougher for a certain group, as well as changes in partially experience-rated employer contributions.

The main findings can be summarized as follows. 1) For older employees a transition to sick leave is often a one-way street out of employment, leading eventually to disability retirement. Half of 50–55 year-olds and over two-thirds of older workers on sickness benefits end up in disability retirement within the next three years. This highlights the importance of preventive measures aimed at minimizing the flow into sick leave. 2) Those employees who could apply for a disability pension under more lenient medical requirements were much more likely to enter sick leave and to retire via disability pension benefits. Therefore, the abolition of the individual early retirement scheme in 2000 significantly reduced the flow into disability retirement in the affected groups. 3) I find strong evidence that experience-rating lowers the flow into sick leave and reduces transitions from sick leave to disability retirement. Moreover, those large firms that can easily bear their share of early retirement costs owing to their strong financial position more readily let employees who are already on sickness benefits exit via disability pension schemes than firms in a weaker financial position. Financial strength does not matter for smaller firms that are not subject to experience-rating. 4) The transition rates to sick leave and disability retirement are relatively large in establishments experiencing a high degree of excess worker turnover. When an establishment is growing, transitions to sick leave and disability retirement become less frequent. There is no evidence that employers exploited the disability pension scheme as a way of adjusting their workforce when downsizing.

These findings imply two policy recommendations to reduce the disability benefit enrolment rate of older workers. First, the stringency of medical criteria and medical screening for disability benefit eligibility should be tough enough.

When non-medical factors are weighted at the expense of medical criteria, disability benefits may distort labour supply decisions, thereby also inducing workers who are not truly disabled to retire via disability programmes. This appears to be mainly a labour supply issue, as I did not find evidence that employers encouraged disability retirement when downsizing. Secondly, the experience-rating of disability benefit costs seems to be an effective policy instrument. It probably induces employers to take preventive action to reduce the inflow into sick leave. Firms also put more effort into getting employees on sickness benefits back into work. This finding should be of considerable interest, not only for Finland, but also for the other countries that do not yet have an experience-rating system for disability benefits. Obviously, there are still a number of open questions regarding, for example, the optimal design of experience-rating and possible spillover effects on hiring and transitions out of work to other destinations than disability retirement.

1.4 Employment and wage effects of payroll tax cut – evidence from a regional experiment⁴

In this paper, I evaluate the employment and wage effects of a regional tax cut experiment in northern Finland. The experiment started in 2003 and it was due to continue for three years, but it was soon extended to 2009 and then again to the end of 2012. The experiment abolished employer contributions to the national pension insurance and national health insurance schemes for firms located in certain high-unemployment regions. Prior to 2003, these employer contributions varied between 2.95 and 6 per cent of the wage bill, depending on the capital intensity and size of the firm. The average payroll tax reduction of the experiment was 4.1 percentage points.

The evaluation setup was designed well before the start of the experiment. First, I chose a comparison region closely resembling the target region and then continued by matching the target region firms to the comparison region firms to form groups of firms as comparable as possible. I did this step at an early stage to make the evaluation transparent and credible. The timing was not very fortuitous, however. The Kainuu self-government experiment, featuring a similar reduction in payroll taxes, started only two years after the experiment in Lapland – and its target area was in the middle of the comparison region. Hence this study considers only the first two years of the experiment.

The main result was that the tax cut did not have a statistically significant effect on employment or the firms' wage sum. The wage sum seemed to have risen but

⁴ Korkeamäki, O. – Uusitalo, R. (2009): Employment and wage effects of a payroll tax cut – evidence from a regional experiment. *International Tax and Public Finance*, 16(6), pp. 753–772.

the standard errors of the estimates were too large to warrant any strong conclusions. These results were based on firm data acquired directly from the tax authorities. I also had data on individual wages from the employer associations. If the estimates from the sub-sample of firms for which wage data is available can be generalized to all firms, about half of the effect of the payroll tax reduction on labour costs was offset by faster wage growth. The remaining two per cent decrease in labour costs did not have a significant effect on employment. According to my estimates, the demand and supply elasticities are roughly equal. The point estimate that the tax cut increased employment by 1.3 percent indicates labour demand elasticity of around 0.6. This is well within the range of earlier estimates. Unfortunately the confidence bands around this estimate are too wide to give much guidance for future tax policy.

The results are in line with findings from other micro-level empirical studies. Usually no employment effects are found and if any effect exists, it is a partial shifting of the tax cut to wages. The most relevant comparisons are with Sweden, where two payroll tax cut experiments were conducted. The first started in 1982 in the four northernmost municipalities and was eventually extended to cover the whole of regional support area A, i.e. almost the northern half of the country, excluding the coast. Payroll taxes were cut by ten percentage points and there was no ceiling to the cut. The experiment, which had become a semi-permanent regional subsidy, had to be phased out by the end of 1999, owing to EU regulations. The second experiment, started in 2002, was similar in geographical scope. It also reduced taxes by ten percentage points but the maximum deductible amount was rather low. Neither of the experiments yielded a statistically significant employment effect. The effect on wages was not investigated for the first experiment and it was positive and significant in the latter case.

The fact that the cut in payroll taxes was targeted at narrowly defined regions and the temporary nature of the tax cut limit the extent to which the results can be generalized to the potential effects of a permanent country-wide reduction in payroll taxes. First, the payroll tax experiment was financed by increasing payroll taxes in the rest of the country. In a national scheme, the budgetary cost would need to be financed by raising other taxes. Second, a regional experiment may have substitution effects if firms reallocate labour to the target region from the rest of the country. This might be beneficial in the sense that part of the reasons for the regional payroll tax cut was to boost employment in disadvantaged regions. However, this limits the usefulness of the results from the experiment in predicting the effects of a national programme. Third, the incidence of the tax cut may also be different in a regional programme since wage contracts are negotiated at the national level. Any nationwide changes in payroll taxes may have an impact on the outcome of these negotiations, while a regional programme that only affects a small share of employers has little weight in national bargaining. Finally, a temporary programme is likely to create smaller employment effects than a permanent reduction in payroll taxes. The expected

duration of three years might not have been a sufficiently long period for firms to adjust their labour demand to a relatively small change in labour costs.

1.5 The Finnish payroll tax cut experiment revisited, or where did the money go?⁵

The payroll tax exemption was originally planned to last for three years, from January 2003 to December 2005. Already in May 2003, the government had decided to start a regional self-government experiment in Kainuu, eastern Finland, beginning from 2005. That experiment contained a similar provision for lowered payroll taxes as the Lapland experiment. Hence it spatially enlarged and temporally extended the payroll tax experiment until the end of 2009. The experiment has been further continued until the end of 2012 and there is intense lobbying to make it a permanent arrangement.

The first evaluation study was somewhat stunted by the start of the Kainuu experiment. The results from the first two years were in the expected direction but less than satisfying in their precision. For this study I had both more firms in the treatment group and more years of observations. I also had information on firms' balance sheets and financial statements that allowed me to track whether the tax cut had an effect on firms' profits, if not on employment or wages. To sum up the situation at the start of writing the fifth essay and to explain the need to re-evaluate the effects of the payroll tax cut. I stated the following:

1) According to our previous research on the first two years of the payroll tax experiment, the tax cut in northern Finland does not seem to have had an immediate employment effect. This finding is consistent with evidence from other Nordic labour markets. 2) In the earlier study there was some indication of rising wages, but not 1:1 with respect to the tax break – this is also a common finding from Sweden and Norway. From 1) and 2) and supported by results from the UK, where a minimum wage change had a negative effect on profits (but no employment effects), it seems likely that changes in payroll tax could also have an effect on firm profitability. Models of incomplete competition from the industrial organisation literature and matching models from the labour market side can accommodate these profit effects, but their size remains an empirical question.

With the larger target area for the experiment, the evaluation setup of the previous study was no longer valid – the Kainuu region formed the main part of the *comparison* region. The target regions in northern Lapland and Kainuu are not geographically linked, but both are within the region eligible for the highest

⁵ Korkeamäki, O. (2011): The Finnish payroll tax cut experiment revisited, or where did the money go? VATT Working Papers 22, Government Institute for Economic Research.

national firm subsidies. From 2000 to 2007, Kainuu was in the highest subsidy region and northern Lapland was part of the region eligible for the second-highest subsidies. There was, however, a special provision that granted firms in Lapland access to subsidies almost as high as for firms in the first category. Instead of doing another matching exercise, I used firms within the same subsidy region but outside the target region of the payroll tax cut as a comparison group (see Figure 6.1).

The target group consists of approximately 2900 firms that in 2001 employed an average of 3.4 workers and had an average turnover of 466,000 euros. The firms in the comparison area were slightly larger, having 3.7 employees and a turnover of 497,000 euros. Prior to the experiment the comparison area firms had also grown somewhat faster than the experiment region firms, but none of the firm-level pre-experiment response measures (levels or growth) differed statistically significantly between the firm groups, even before controlling for industry, growth trends, *etc.* Therefore I claim that the setup was rather successful. The analysis was done with differences-in-differences regressions, with controls for either industry- and region-specific or firm-specific growth trends.

The main result from this new study is that the payroll tax cut did not have a statistically significant effect on employment, the wage sum, profits or hourly wages in the private sector. Most of the estimates are positive but unfortunately the standard errors are so wide that they could accommodate values indicating a full shifting of the tax cut to either the wage sum, profits or, indeed, to employment. If we look at the point estimates in euro or employee terms, they indicate that the wage sum in the target region firms rose faster than in the comparison region, employment growth did not react and profits grew even a little more than the wage sum. Alternatively, if we consider the point estimates of the percentage changes, employment and the wage sum did grow by an equal amount and there was no effect on profits. The only unambiguous finding is that the tax cut is reflected in the financial statement data, although, even there, there was some uncertainty in the case of small and the least capital-intensive firms.

The effect on hourly wages found in our earlier study is not found for the combined target region of Kainuu and Lapland. The effect is still found for Lapland – but the estimates for Kainuu would imply a negative wage effect and the effects cancel each other out when calculating the total. The results also show one statistically significant change in a non-experiment year that calls for caution in interpreting the results. Certainly, a region-specific shock could have taken place in Kainuu and caused the negative effect, but I found no reason to believe that the result for Lapland is trustworthy.

Irrespective of the findings of this and other similar experiments in other Nordic countries, national pension insurance payment contributions have gradually been lowered in recent years on the premise that this is a cost-effective way of

boosting employment. From the beginning of 2010, they were abolished altogether. There was some debate as to whether this was the most effective way to help firms to generate jobs, but empirical findings had a rather minor role in the discussion. This is, of course, partly due to the lack of conclusive findings.

The Finnish payroll tax experiment is a rare example of a tax change being made in an experimental setting with the stated aim of facilitating economic research. Hence it was important to evaluate the experiment, even if the results tell rather little about the effects. This was also an opportunity to learn about the experiment itself in order to understand better how possible future experiments should be designed and implemented to the greatest scientific advantage. I argue that it is still important to continue experimenting – it is also important to pre-evaluate experiments to see if they are likely to yield accurate and reliable results.

2. A Gender wage gap decomposition for matched employer-employee data⁶

Abstract

In this paper, we evaluate the extent to which the gender wage gap in the Finnish manufacturing sector is attributable to within-job wage differentials, gender differences in individual qualifications, and a disproportionate concentration of women in lower-paying firms and lower-paying jobs within firms. We use matched employer-employee data to compare wage differentials between similarly qualified female and male workers who are doing similar work for the same employer. Our modelling approach employs a correlated random effects specification to account for the hierarchical grouped structure of the underlying data.

Key words: Gender wage gap, wage discrimination, gender segregation, random effects model

JEL classification numbers: J14, J23, J26

2.1 Introduction

A huge body of literature has emerged to explain why the gender wage gap persistently exists in virtually all labour markets (see Altonji and Blank, 1999, and Blau and Kahn, 2000, for recent surveys). Traditional attempts to explain the wage gap focused on gender differences in individual qualifications and their rewards in the labour market. More recently, the importance of the segregation of women and men into different jobs has been recognized. This line of research emphasizes that wages are closely tied to the characteristics of jobs, not only to the individuals who hold them. If typical female jobs pay lower wages than jobs dominated by men, the mean earnings of women can fall short of men's earnings even in the absence of within-job wage differentials between the sexes.

⁶ Ossi Korkeamäki and Tomi Kyyrä.

The earlier version of the paper was circulated under the title "Explaining gender wage differentials: Findings from a random effects model". We appreciate the helpful comments received at the third Nordic Workshop on the Economic Analysis of Linked Employer-Employee Data in Bergen, the second Nordic Econometrics Meeting in Bergen, the CAED Conference in London, and the EALE Conference in Seville. We are grateful to the Confederation of Finnish Industry and Employers for access to their data. Antti Luukkonen kindly provided his measures of job complexity levels for our use. The suggestions of a co-editor and anonymous referees considerably improved the paper.

Attempts to quantify the segregation effects on the wage gap were for a long time distorted by a lack of appropriate data. Consequently, most early analyses focused on segregation among occupations, firms, or industries only. This is clearly unsatisfactory, as women and men are further segregated into different jobs within firms. In recent years, important advances have been made by access to large matched employer-employee data sets that contain multiple observations on workers with the same employer. When information on occupations or job titles is available, such data enable wage comparisons between male and female workers who are doing similar work for the same employer. This kind of comparative analysis has been conducted by Petersen and Morgan (1995), Petersen *et al.* (1997), Meyersson Milgrom *et al.* (2001), Groshen (1991), Datta Gupta and Rothstein (2001), and Bayard *et al.* (2003). In the first three of these studies observed sex differentials in mean wages within jobs are simply aggregated to form various wage decompositions. This approach has the obvious drawback that variation in individual characteristics is left uncontrolled. In the other studies, wages are regressed against a set of control variables and fraction female in the worker's industry, firm, occupation, and/or job.⁷ The key idea is that the regression coefficients of the various fraction female variables capture the relationship between the wage rate and 'femaleness' of the underlying labour market structure.

It should be noted that a common practice in the fraction female regressions above has been to neglect the grouping in the underlying data. For example, observations on workers resulting from the same firm are interpreted as being independent.⁸ However, intuition suggests that we should expect workers in the same firm to be more homogeneous than those in a sample drawn randomly from the population of all firms. Workers in the same firm share many common factors, some of which may be observable (e.g. firm size, fraction female) but many are not (e.g. market power, managerial ability). In the regression analysis the effect of such unobservables serves as a latent firm effect that will be absorbed into the error term. Moreover, since different jobs require different skills and qualifications, we can further expect that within a given firm workers who are doing the same job are more homogenous than the firm's workforce as a whole. This implies an additional source of dependence between workers within jobs.

In general, the matched employer-employee data exhibit a particular type of grouped structure, which contrasts the statistical properties of such data with the classical random sample case. A consequence of the grouping in the regression

⁷ In a related paper, we apply this method to the Finnish data; see Korkeamäki and Kyrrä (2002). Groshen's (1991) specifications do not include control variables.

⁸ Bayard *et al.* (2003) report the standard errors adjusted for intra-establishment error correlation.

analysis is that, owing to the latent group effects, the errors will be correlated within groups. In the absence of correlation between the latent group effects and regressors included in the model, the OLS coefficients will be consistent, but the usual standard error estimates can be very misleading (Moulton, 1986). More generally, when the group effects are correlated with the regressors, the OLS coefficients will be inconsistent.

In this study, we explore wage differentials between white-collar women and men in the Finnish manufacturing sector using a large matched employer-employee data set. We view the data as having a nested structure with three levels: firms, jobs within firms, and workers in jobs within firms.⁹ A job is defined as an occupation within a firm. Along with individual characteristics, the wage rate is allowed to depend upon the employing firm and the job the worker holds within the firm. The latent firm and job effects are modelled as functions of group characteristics, including the mean characteristics of individuals within the groups. We end up with a regression model with variables measured at the individual, job, and firm levels, and an error term that has a two-way nested structure with separate intercepts for firms and jobs within firms. Using the regression results we decompose the overall gender gap in pay into the contributions of gender segregation, gender differences in the individual qualifications, and the unexplained within-job gap.

Our approach departs from the existing segregation literature in that we explicitly model wage differentials between firms and jobs. In contrast to standard fraction female regressions, we obtain consistent estimates of the parameters of interest in the presence of the correlated group effects that are likely to arise in the case of the matched employer-employee data. With respect to job segregation, the previous studies have focused on quantifying what fraction of the overall wage gap can be attributed to a disproportionate concentration of women in lower-paying jobs. In addition to identifying this quantity under less restrictive assumptions, we go a step further by addressing the issue of why typical female jobs are lower-paid. When evaluating the extent to which lower wages in predominantly female jobs can be explained by job attributes, we make use of an index of job complexity that measures the responsibility, skills and effort required by a given job. Thus we are able to assess whether wage differentials between typical female and male jobs can be viewed as justified or not, a question that is beyond the scope of earlier analysis but crucial, for example, in the view of comparable worth policy.

⁹ Obviously, we could go further and introduce an additional level on top of this hierarchy by grouping firms by industry. For simplicity, we focus on the three levels and treat industry as a characteristic of firms rather than a hierarchy level of its own.

In the next section, we describe the econometric methods and contrast our approach with the fraction female decompositions in the previous studies. Section 2.3 gives details about the data and reports some descriptive statistics. The results are reported in Section 2.4, which is followed by a concluding section.

2.2 Methodological framework

2.2.1 The wage model

Suppose our data consist of all employees of F firms. Within firms employees who do similar work are grouped together, in which case they are said to hold the same job. Observations across firms are regarded as being independent, but within firms wages are correlated owing to common firm and job characteristics. We model the log wage of worker i ($i = 1, 2, \dots, n_{jk}$) who holds job k ($k = 1, 2, \dots, c_j$) in firm j ($j = 1, 2, \dots, F$) as

$$w_{jki} = \eta s_{jki} + \boldsymbol{\beta}' \mathbf{x}_{jki} + f_j + v_{jk} + \varepsilon_{jki}, \quad (1)$$

where s is the female dummy, \mathbf{x} is a vector of other individual characteristics, v is the job effect that is nested within the firm effect f . For the idiosyncratic errors ε , we assume

$$E(\varepsilon_{jki} | \mathbf{X}_j, \mathbf{v}_j, f_j) = 0 \text{ and } E(\varepsilon_{jki} \varepsilon_{jki'} | \mathbf{X}_j, \mathbf{v}_j, f_j) = 0 \text{ for } i \neq i', \quad (2)$$

where \mathbf{X}_j includes \mathbf{x} and s for all employees of firm j , and $\mathbf{v}_j = (v_{j1}, v_{j2}, \dots, v_{jc_j})$.

Wage variation between firms and jobs beyond the observable individual characteristics is captured by f and v respectively. Without loss of generality, the job effects are defined in deviation from the firm effects, with the expected value within each firm equal to zero. Thus, $E(f + v | \text{firm } j) = f_j + E(v | \text{firm } j) = f_j$. We emphasize that f and v are likely to be correlated with s and \mathbf{x} . In particular, women are expected to be concentrated in firms with low values of f , and further in jobs with low values of v . Since different firms and jobs require different qualifications, the group effects f and v are likely to be correlated also with the variables in \mathbf{x} . If $f_j > f_{j'}$, workers in firm j earn more on average than workers in firm j' after controlling for s and \mathbf{x} . Similarly, provided that $v_{jk} > v_{jk'}$, workers in job k are more highly paid on average than those in job k' within the same firm j , after controlling for s and \mathbf{x} .

Within jobs wage differentials are related to workers' sex (s), other individual characteristics (\mathbf{x}), and unobservables (ε). A parameter of particular interest is η that gives the expected wage differential between equally qualified (in terms of \mathbf{x}) women and men who are doing the same work for the same employer. One

may be tempted to view a negative value of η as evidence of wage discrimination against women. Such an interpretation is justified only if all relevant explanatory variables were included in \mathbf{x} . This may not be the case in practice. In general, the influences of possible discrimination and unmeasured individual characteristics are indistinguishable in the value of η . Therefore, we interpret η simply as a measure of the unexplained within-job wage differential between sexes.

At this point, a few remarks on the restrictions imposed above are in order. First, the returns to individual qualifications, $\boldsymbol{\beta}$, are assumed equal for women and men. One should recognise that the interpretation of $\boldsymbol{\beta}$ is conditional on the position held in the labour market (i.e. conditional on f and v), so $\boldsymbol{\beta}$ measures the returns within a given job. Since employers cannot apply very different reward schemes to their female and male employees who are doing the same work, our assumption is not as restrictive as it might first look. We will return to this issue and present results from a regression model with gender-specific slopes. The assumption that the unexplained within-job wage gap is of the same size everywhere is rather restrictive. One might wish to allow the coefficient of the female dummy to vary across jobs, i.e. replace η with η_{jk} . We adopt a very narrow definition for jobs in our empirical application. This results in a huge number of jobs, many of which include either female or male employees only, making the estimation of job-specific coefficients infeasible in practice.

2.2.2 Decomposing the gender gap in pay

The gender wage gap is defined as the difference in the expected wages between men and women, i.e. the wage difference between a randomly chosen man and woman. Using the model outlined above we decompose it as

$$\begin{aligned} E(w | s = 0) - E(w | s = 1) = & -\eta + \boldsymbol{\beta}' [E(\mathbf{x} | s = 0) - E(\mathbf{x} | s = 1)] \\ & + [E(f | s = 0) - E(f | s = 1)] \\ & + [E(v | s = 0) - E(v | s = 1)], \end{aligned} \quad (3)$$

where the contributions of gender segregation among firms and jobs are captured by the last two terms. A positive value of $E(f | s = 0) - E(f | s = 1)$ indicates that women are disproportionately concentrated in lower-paying firms. This term would be zero, if there were no variation in f across firms or if women and men were identically distributed across firms. If women are relatively more frequently allocated to lower-paying jobs within firms, $E(v | s = 0) - E(v | s = 1)$ will take a positive value. It would be zero, if there was no systematic wage variation across jobs within firms beyond the differences in individual characteristics or if, within all firms, women and men were allocated identically across jobs. The amount of within-job wage differentials between sexes not accounted for by the explanatory variables \mathbf{x} equals $-\eta$. The contribution of gender differences in individual characteristics is captured by the remaining term on the right-hand side.

To obtain an empirical counterpart of the decomposition, the conditional means of w and \mathbf{x} can be replaced with the sample means over women and men but the other components need to be estimated. Since the latent group effects are expected to be correlated with the explanatory variables, we will focus on estimation by fixed effects and correlated random effects.

2.2.3 The fixed effects approach

In our first approach, we take f and v as fixed constants to be estimated along with η and $\boldsymbol{\beta}$. Therefore, we consider the model conditional on the firm and job effects:

$$E(\bar{w}_{jki} | \mathbf{X}_j, \mathbf{v}_j, f_j) = \eta s_{jki} + \boldsymbol{\beta}' \mathbf{x}_{jki} + f_j + v_{jk}. \quad (4)$$

In this case, η and $\boldsymbol{\beta}$ could be estimated by regressing w on s , \mathbf{x} and the full set of job dummies. As the number of job dummies may be too large to make estimation feasible, we obtain analytically equivalent estimators of η and $\boldsymbol{\beta}$ by applying pooled OLS to the transformed model:

$$w_{jki} - \bar{w}_{jk.} = \eta (s_{jki} - \bar{s}_{jk.}) + \boldsymbol{\beta}' (\mathbf{x}_{jki} - \bar{\mathbf{x}}_{jk.}) + \varepsilon_{jki} - \bar{\varepsilon}_{jk.}, \quad (5)$$

where $\bar{w}_{jk.}$, $\bar{s}_{jk.}$, $\bar{\mathbf{x}}_{jk.}$, and $\bar{\varepsilon}_{jk.}$ denote averages over workers in the k -th job of firm j . Under the assumptions (2), the resulting ‘‘fixed effects’’ (FE) estimators $\hat{\eta}$ and $\hat{\boldsymbol{\beta}}$ are consistent under arbitrary correlation between (s, \mathbf{x}) and (f, v) . Given the restriction $E(v | \text{firm } j) = 0$ for all j , the firm and job effects can be estimated as

$$\hat{f}_j = \bar{w}_{j..} - \hat{\eta} \bar{s}_{j..} - \hat{\boldsymbol{\beta}}' \bar{\mathbf{x}}_{j..}, \quad (6)$$

$$\hat{v}_{jk} = \bar{w}_{jk.} - \hat{\eta} \bar{s}_{jk.} - \hat{\boldsymbol{\beta}}' \bar{\mathbf{x}}_{jk.} - \hat{f}_j, \quad (7)$$

where $\bar{w}_{j..}$, $\bar{s}_{j..}$ and $\bar{\mathbf{x}}_{j..}$ denote averages over the employees of firm j . The point estimates of f and v are noisy because the number of observations per firm and, especially, per job can be small. However, the estimates of their expected values among women and men based upon sample averages are expected to be reasonably accurate. Thus, we proceed by inserting $\hat{\eta}$ and $\hat{\boldsymbol{\beta}}$ along with the sample means of \hat{f} and \hat{v} over women and men into (3). This gives the first version of our wage gap decomposition. It allows us to distinguish the contributions of gender segregation among firms and jobs to the overall wage gap from the contributions of the unexplained within-job gap and gender differences in individual characteristics.

2.2.4 The correlated random effects approach

In an alternative approach, we take an explicit account of the relationship between the latent group effects and the explanatory variables. More precisely, we specify the expected values of f and v conditional on observables via auxiliary linear regressions. Let $\mathbf{X}_j^* = (\mathbf{X}_j, \mathbf{z}_j, \mathbf{g}_{j1}, \dots, \mathbf{g}_{jc_j})$ be the extended set of conditioning variables that includes firm attributes \mathbf{z}_j (firm size, industry, etc.) and job attributes \mathbf{g}_{jk} 's (job size, job complexity index, etc.) in addition to \mathbf{X}_j . We specify the conditional mean of the firm effect as

$$E(f_j | \mathbf{X}_j^*) = \alpha + \delta_0 \bar{s}_{j..} + \delta_1' \bar{\mathbf{x}}_{j..} + \delta_2' \mathbf{z}_j \quad (8)$$

and that of the job effect as

$$E(v_{jk} | \mathbf{X}_j^*) = \theta_0 (\bar{s}_{j.k.} - \bar{s}_{j..}) + \theta_1' (\bar{\mathbf{x}}_{j.k.} - \bar{\mathbf{x}}_{j..}) + \theta_2' (\mathbf{g}_{jk} - \bar{\mathbf{g}}_{j.}), \quad (9)$$

i.e. the first moments of the marginal distributions of f and v are assumed to be linear functions of the group means of s and \mathbf{x} and of other group level variables. All the explanatory variables on the right-hand side of (9) are measured in deviation from the firm mean in order to enforce the expected value of v within firms to zero.¹⁰

Now we consider the model conditional on \mathbf{X}_j^* :

$$E(w_{jki} | \mathbf{X}_j^*) = \eta s_{jki} + \beta' \mathbf{x}_{jki} + E(f_j | \mathbf{X}_j^*) + E(v_{jk} | \mathbf{X}_j^*).$$

Defining $\zeta_j \equiv f_j - E(f | \mathbf{X}_j^*)$ and $\omega_{jk} \equiv v_{jk} - E(v_{jk} | \mathbf{X}_j^*)$, we obtain the estimating wage equation:

$$\begin{aligned} w_{jki} = & \alpha + \eta s_{jki} + \delta_0 \bar{s}_{j..} + \theta_0 (\bar{s}_{j.k.} - \bar{s}_{j..}) + \beta' \mathbf{x}_{jki} + \delta_1' \bar{\mathbf{x}}_{j..} + \theta_1' (\bar{\mathbf{x}}_{j.k.} - \bar{\mathbf{x}}_{j..}) \\ & + \delta_2' \mathbf{z}_j + \theta_2' (\mathbf{g}_{jk} - \bar{\mathbf{g}}_{j.}) + u_{jki}, \end{aligned} \quad (10)$$

where $u_{jki} \equiv \zeta_j + \omega_{jk} + \varepsilon_{jki}$. Conditional on \mathbf{X}_j^* , all components of u_{jki} are assumed to be mutually independent, with zero means and constant variances σ_ζ^2 , σ_ω^2 , and σ_ε^2 respectively. Within firm j , the variance-covariance structure of the errors is given by

¹⁰ This is only a matter of parameterization provided that $\bar{\mathbf{g}}_{j.}$ is included in the set of firm covariates \mathbf{z}_j .

$$E(u_{jki}u_{jk'i'} | \mathbf{X}_j^*) = \begin{cases} \sigma_\xi^2 + \sigma_\omega^2 + \sigma_\varepsilon^2, & \text{if } k = k' \text{ and } i = i'; \\ \sigma_\xi^2 + \sigma_\omega^2, & \text{if } k = k' \text{ and } i \neq i'; \\ \sigma_\xi^2, & \text{if } k \neq k' \text{ and } i \neq i'. \end{cases} \quad (11)$$

This is known as the two-way nested error structure in econometrics (Fuller and Battese, 1973). It models the residual correlation within firms that remains after conditioning on the observed firm, job, and individual characteristics. Such a correlation is likely to exist owing to unobservable job and firm factors. We estimate the model with generalized least squares (GLS) that exploits the particular form of the error structure for efficiency and produces appropriate standard errors.¹¹ It should be stressed that including the group means of individual explanatory variables in (8) and (9) provides a way of allowing s and \mathbf{x} to be correlated with f and v , an old idea by Mundlak (1978).¹² To emphasize this point, we refer to the specification outlined above as the “correlated random effects” (CRE) model.¹³

Coefficients of the fraction female variables in (8) and (9) are of particular interest. A negative value of δ_0 implies that firms with a high density of female workers pay lower wages after controlling for $\bar{\mathbf{x}}_{j..}$ and \mathbf{z}_j . If within firms employees in predominantly female jobs are lower paid given $(\bar{\mathbf{x}}_{jk.} - \bar{\mathbf{x}}_{j..})$ and $(\mathbf{g}_{jk} - \bar{\mathbf{g}}_j)$, it will be indicated by a negative value of θ_0 . In other words, δ_0 and θ_0 are kind of “residual gender effects”, which imply that predominantly female firms and jobs pay different wages for reasons not accounted for by the observed worker and group characteristics.

Because $E(f | s = 0) = E[E(f | \mathbf{X}^*) | s = 0]$ by the law of iterative expectations, we obtain an estimate of $E(f | s = 0)$ by averaging the right-hand side of (8) over all men. $E(f | s = 1)$ is estimated analogously by averaging over women. The contribution of gender segregation among firms can then be expressed as

¹¹ The large unbalanced data raise some computational issues, as the inverse of the error variance-covariance matrix is required by the GLS procedure. These issues and the estimation of the variance components are discussed in Korkeamäki and Kyyrä (2003).

¹² Chamberlain (1984) considers a general case where the latent group effects are modelled as linear predictors of s and \mathbf{x} of all employees within the group. Mundlak’s (1978) specification is obtained by imposing a restriction that the coefficients of s and \mathbf{x} in the linear predictor are identical for all i within the group. The unrestricted specification becomes cumbersome in our case where group sizes vary and some firms are very large.

¹³ The model defined by (10) and (12) is known under a variety of other names, including the nested error components model, variance components model, random intercepts model, mixed model, and hierarchical model.

$$\begin{aligned}
E(f | s = 0) - E(f | s = 1) &= \sum_{j=1}^F (o_j^m - o_j^f) \delta_0 \bar{s}_{j..} \\
&+ \sum_{j=1}^F (o_j^m - o_j^f) (\delta_1' \bar{\mathbf{x}}_{j..} + \delta_2' \mathbf{z}_j),
\end{aligned} \tag{12}$$

where o_j^f (o_j^m) is the fraction of all women (men) allocated to firm j . Similarly, the contribution of gender segregation among jobs within firms is given by

$$\begin{aligned}
E(v | s = 0) - E(v | s = 1) &= \sum_{j=1}^F \sum_{k=1}^{c_j} (o_{jk}^m - o_{jk}^f) \theta_0 (\bar{s}_{jk.} - \bar{s}_{j..}) \\
&+ \sum_{j=1}^F \sum_{k=1}^{c_j} (o_{jk}^m - o_{jk}^f) \left[\theta_1' (\bar{\mathbf{x}}_{jk.} - \bar{\mathbf{x}}_{j..}) + \theta_2' (\mathbf{g}_{jk} - \bar{\mathbf{g}}_{j..}) \right],
\end{aligned} \tag{13}$$

where o_{jk}^f (o_{jk}^m) is the fraction of all women (men) allocated to the k -th job of firm i . Substituting (12) and (13) into (3) along with the GLS estimates of the regression coefficients gives us the second version of our wage gap decomposition.

In the case of the CRE model, the segregation contributions can be expressed as sums of various terms. These terms pass on useful information, which is not available from the FE model. For example, if typical female jobs are found to be characterised by low values of v , one may speculate that lower wages in such jobs result from lower skill requirements. If this is the case, a large fraction of the contribution of gender segregation among jobs in (13) will be attributed to differences in the mean education (incorporated in $\bar{\mathbf{x}}_{jk.} - \bar{\mathbf{x}}_{j..}$) and job complexity (incorporated in $\mathbf{g}_{jk} - \bar{\mathbf{g}}_{j..}$), while the component associated with the fraction female ($\bar{s}_{jk.} - \bar{s}_{j..}$) will be close to zero. By contrast, if wage differentials between typical female and male jobs arise largely from some unobserved sources, this will be indicated by a strong effect of the fraction female term in (13).

2.2.5 Discussion

In the case of the FE model, we cannot say anything about why predominantly female firms and jobs are lower paid on average. This of course is a cost of the robustness of the fixed effect method: we do not assume anything about the relationship between the group effects and regressors. Compared with the FE model, the CRE specification is more restrictive, as the conditional expectations of f and v are assumed linear. However, when the group means of s and \mathbf{x} are included in (8) and (9), the GLS estimators of η and $\boldsymbol{\beta}$ are identical to their FE estimators, and hence not affected by these additional restrictions. In this respect, we do not lose anything by imposing more structure on the model. The additional

structure of the CRE model is exploited in explaining wage variation between firms and jobs. While both the FE and CRE model are able to produce identical results for the effects of individual-level regressors, only the latter is informative about wage differentials between firms and jobs. For this reason, the CRE model is our preferred choice. A potential drawback of the method is that the sum of various contributions does not necessarily equal the raw wage gap in finite samples. This is a consequence of the more complex error structure.

Our approach departs from the decomposition exercises in Groshen (1991), Datta Gupta and Rothstein (2001), and Bayard *et al.* (2003) in some essential ways. Of course, the main difference is that our CRE approach is informative about the determinants of lower wages in predominantly female firms and jobs. Secondly, the interpretation of the regressor coefficients η and β comes from the wage model defined in (1) and (2), i.e. they measure wage differentials within jobs (In other words, the firm and job effects held constant). This interpretation is trivial when the model is estimated by fixed effects, but the coefficients have the same meaning also in the CRE specification as we allow f and v to be correlated with s and x . The coefficients in the standard fraction female regressions do not generally have the same interpretation. Thirdly, we define the segregation contributions as differences in the mean values of the firm and job effects between men and women.¹⁴

Despite the differences in the modelling framework, the estimating wage equation in (10) and the associated decomposition are not much different from those in the previous studies. If occupational segregation is omitted, the standard fraction female decompositions can be viewed as special cases of our CRE decomposition. If we set δ_1 , δ_2 , θ_1 , and θ_2 to zero, we obtain a specification similar to those in Datta Gupta and Rothstein (2001) and Bayard *et al.* (2003). If we impose further $\beta = \mathbf{0}$, the model is reduced to Groshen's (1991) specification. Within our framework, the restriction $\delta_1 = \theta_1 = 0$ is equivalent to assuming that the firm and job effects are uncorrelated with x . This of course is a rather restrictive assumption, and it may lead to inconsistent estimates of η and β . The importance of this sort of restrictions is an empirical issue, and it depends on the data in hand. For example, both Datta Gupta and Rothstein (2001) and Bayard *et al.* (2003) find only a minor change in the female dummy coefficient when the fraction female variables were replaced with the full set of job dummies. In general, it does make a difference whether one conditions on the job held or only on the femaleness of the worker's position. In our application, we find quantitatively significant discrepancies in the estimated coefficients, standard

¹⁴ Additional, less important, differences are: (1) we measure the fraction female in job as a deviation from the firm mean, (2) we do not include the fraction female in occupation nor in industry in our model, and (3) we apply GLS, not OLS.

errors, and decomposition results between the CRE model and standard fraction female specification.

2.3 Data and descriptive statistics

2.3.1 The TT data and job classification

Our data come from the records of the Confederation of Finnish Industry and Employers (TT). TT is the central organisation of manufacturing employers and its member firms account for more than three-quarters of the value added of the Finnish manufacturing sector. Each year TT conducts three surveys covering almost all employees of its member firms. All surveys are directed to the employer, one asking information about white-collar workers and the other two about blue-collar workers. The focus of this paper is restricted to white-collar workers because of differences in the available records and compensation schemes between the two worker groups.¹⁵ In the subsequent analysis, we include all the full-time white-collar workers aged between 18 and 65 who were employed in 2000 by TT firms with at least 5 workers. The resulting data contain observations on almost 150,000 employees in 1,464 firms. This is not a representative sample of the Finnish manufacturing firms: while all the large firms are included, many small and medium sized firms are not members of TT and hence do not appear in the data. However, the data are rather representative of the white-collar population, covering roughly 60% of the salaried employees in the Finnish manufacturing sector. Compared with the entire population of white-collar manufacturing workers, individuals in the data are somewhat more educated and slightly higher paid, but there are no notable differences in the sex or age composition, or in the allocation of workers across different industries.¹⁶

In Finland working conditions are regulated by collective agreements, which are made along industrial lines between employer organisations and unions. In the collective agreements, white-collar workers are grouped into managerial, technical, and clerical workers, where the latter two groups are combined in some industries. Each group is covered by a separate industry-specific agreement, which determines, among other working conditions, a minimum rate of pay for a particular type of work. For this purpose, job tasks of technical and clerical employees are evaluated according to the responsibility, skills, and effort they

¹⁵ The working paper version includes a separate analysis for blue-collar workers; see Korkeamäki and Kyyrä (2003).

¹⁶ Of course, white-collar manufacturing workers as a whole are a very selective group of all workers. Compared with the most other worker groups, women are underrepresented and the gender wage gap is higher within this group. We focus on this narrow group to be able to exploit information on the unique measure of job complexity. In Section 4, we briefly discuss wage gap decomposition results for blue-collar manufacturing workers and workers in the private service sector.

require, and thereby mapped onto a scale of complexity levels. Each level of complexity is then associated with a given basic wage that serves as the wage floor for that type of work. It should be stressed that the employers do not hold the evaluation of job tasks in the palm of their hands; it is highly regulated and supervised by the representatives of unions. The key principle is that the basic wage is determined by job attributes only, independently of the individual characteristics of the jobholder (e.g. sex and education). Actual wages received by workers generally exceed the basic wages, owing to firm premiums and rewards for individual qualifications and performance.

Workers are also classified into 78 occupational groups that are common to all white-collar groups and all industries. We define a job as an occupation within the employing firm (this results in 26,236 jobs).¹⁷ However, where workers covered by different agreements are allocated to the same job, we split the job into parts, each one including only technical, clerical, or managerial employees. The number of jobs increases to 30,281. Finally, jobs that include workers with differing levels of job complexity are further divided into jobs including only workers with the same level of job complexity. This raises the number of jobs to 40,664. At the end, all workers within a given job have the same occupation and job complexity classification and are covered by the same collective agreement.

The evaluation process of jobs explicitly states that wages are closely tied to jobs, not only to the workers who hold them. The knowledge of job complexity ranking provides valuable information that can be used to explain wage variation between jobs. There are two complications, however. First, job complexity information is missing for all the managerial jobs, which are not subject to any evaluation process. Secondly, the scale of the complexity classification is not constant but 9 different scales are applied in different industries. Where no distinction between the technical and clerical employees is made, the number of complexity levels lies somewhere between 3 and 15 (being 8, 9, or 10 in most cases). In other industries the number of complexity levels is 6 for technical employees and 12 for clerical employees. However, the different scales for job complexity cover roughly the same range of logarithmic basic wages and the relationship between the job complexity levels and basic wages is approximately log-linear within each scale. Therefore, we re-scale the original complexity

¹⁷ Additionally, we use information on each worker's job location (municipality) to sub-divide jobs within firms. Workers with the same occupation who are working in different plants of the same firm are allocated to separate jobs, if the plants are located in different municipalities. We do not consider it prudent to divide firms into plants with this indirect information on job locations. Thus, the employer unit in our analysis is firm, not plant.

variables on the interval 0 to 9 by applying a suitable stretch or compression factor to each industry-specific scale.¹⁸

2.3.2 Sample statistics

From Table 2.1 we see that some 37% of white-collar workers are female and they earn on average 22% less than their male counterparts do. There are no large gender differences in the average age, work experience, or firm tenure. Women's slightly longer average experience and tenure reflect the strong labour market attachment of the Finnish women. On the other hand, the potential years of work experience has a tendency to overestimate the actual years spent in the labour market especially for women. Firm tenure is not subject to such a bias, however. Roughly equal mean tenures for women and men suggest that mobility patterns are rather similar for both sexes. This view is in accordance with the results of Lilja (1995) who finds no evidence of gender differences in the exit rates from the TT firms for white-collar women and men. While men are only slightly more educated as measured by the education level, gender differences in terms of the field of education are quite substantial. Of men, 65% have received a technical education, compared with 17% of women. Moreover, 41% of women have obtained a degree in social sciences, business, or law. The mean value of the job complexity level is clearly higher for men, indicating that more demanding clerical and technical jobs are mainly occupied by men.

To give a hint of the role of gender segregation, Table 2.2 shows the gender wage ratio and sex composition by 2-digit occupation group and white-collar group (i.e. managerial vs. technical and clerical workers). Variation in the female share indicates a large degree of gender segregation among occupations, perhaps reflecting differences in education. Women appear to be concentrated in the administrative occupations. By contrast, less than 10% of white-collar workers in production occupations are female. The gender wage ratio within the occupation groups ranges from .670 to 1.062, being on average clearly higher than the raw wage gap on the bottom line. This suggests that occupational segregation plays a role in explaining the gender wage gap. There is no clear relationship between the female share and the size of the within-occupation gender wage gap: the correlation coefficient between these variables is .26 and statistically insignificant at the 5% level.¹⁹

¹⁸ This conversion idea came from Antti Luukkonen, who found by comparing various wage regressions that the single variable works relatively well compared with the huge number of industry-specific complexity dummies.

¹⁹ When the observations are weighted by the size of the occupation group, the correlation coefficient is even lower (.06).

Table 2.1 *Sample statistics by gender*

	Women		Men		All	
Hourly wage, euro	12.107	(3.926)	15.763	(5.59)	14.409	(5.336)
Log hourly wage	2.452	(.276)	2.702	(.326)	2.610	(.331)
Age	41.146	(9.794)	41.128	(9.920)	41.135	(9.874)
Schooling years	12.065	(2.164)	12.810	(2.200)	12.534	(2.216)
Firm tenure, years	12.395	(10.648)	11.860	(10.552)	12.058	(10.591)
Work experience, years	22.081	(10.632)	21.318	(10.381)	21.601	(10.481)
Job complexity (0–9 scale)	3.707	(1.69)	4.909	(2.201)	4.379	2.078
Job complexity missing, share	.437		.569		.520	
Job size	242	(665)	310	(754)	285	(723)
Employer size	3,445	(4,772)	3,526	(4,759)	3,496	(4,764)
Education level, %						
Basic or unknown	19.256		11.966		14.665	
Secondary	29.916		22.621		25.322	
First stage of tertiary	31.544		27.030		28.701	
Bachelor's degree	8.247		22.559		17.259	
Master's degree	10.528		14.808		13.223	
PhD	.509		1.017		.829	
Field of education, %						
General	8.070		6.382		7.007	
Education	.372		.069		.181	
Humanities and art	2.825		0.413		1.306	
Social sciences, business and law	40.975		9.500		21.156	
Science	2.687		2.491		2.563	
Technical	17.091		64.722		47.083	
Agriculture	1.115		3.062		2.341	
Health and welfare	2.591		0.364		1.188	
Services	5.006		1.029		2.502	
Unknown	19.270		11.969		14.673	
Fraction female in firm	.436	(.165)	.332	(.137)	.370	(.156)
Fraction female in job	.778	(.293)	.131	(.185)	.370	(.388)
Sample size	55,158		93,786		148,944	

Notes: Unless otherwise indicated, the figures in the table are means. Standard errors are in parentheses. Hourly wage is computed dividing the monthly wage by regular working hours. Schooling years is defined as the mean years of schooling attached to a given level of education. Work experience is approximated by subtracting the years of schooling and seven years for time prior to the age of school entry from the worker's age. The mean level of job complexity is computed using non-missing values only. Employer and job sizes are the average firm and job size over workers. The mean firm size in the data is 102 and the mean job size is 3.7. The total number of firms is 1,464 and that of jobs is 40,664.

Table 2.2 Fraction female and gender wage ratio by 2-digit occupation and white-collar group

Occupational group	Managerial			Technical & clerical			All white-collar		
	<i>N</i>	Fem	Gap	<i>N</i>	Fem	Gap	<i>N</i>	Fem	Gap
R&D	28,128	.163	.909	19,102	.422	.791	47,230	.268	.759
R&D management	1,320	.135	.847	4,956	.528	.859	6,276	.445	.614
Product design	21,444	.133	.914	8,824	.279	.826	30,268	.176	.843
Quality management	1,397	.327	.894	3,316	.555	.854	4,713	.487	.800
Research	3,967	.277	.885	2,006	.572	.846	5,973	.376	.781
Production	8,469	.079	.842	24,627	.102	.865	33,096	.096	.841
Production and maintenance management	5,498	.053	.864	17,810	.074	.857	23,308	.069	.836
Production support	2,971	.127	.858	6,817	.177	.866	9,788	.162	.842
Logistics	2,082	.230	.812	4,870	.460	.908	6,952	.391	.794
Materials and logistics	472	.138	.825	2,515	.274	.883	2,987	.252	.825
Purchasing	1,467	.241	.811	1,647	.583	.842	3,114	.422	.725
Shipping	143	.413	.844	708	.833	.930	851	.763	.805
Sales and marketing	8,807	.231	.831	14,231	.663	.765	23,038	.498	.662
Sales	7,211	.207	.823	12,720	.681	.767	19,931	.510	.643
Sales promotion	709	.453	.845	721	.544	1.009	1,430	.499	.876
Production and marketing co-operation	887	.246	.828	790	.467	.764	1,677	.350	.766
PR	1,828	.398	.884	2,992	.575	.846	4,820	.508	.803
PR	650	.697	.836	629	.812	.891	1,279	.754	.806
Information technology	1,178	.233	.861	2,363	.512	.817	3,541	.419	.750
Juridical & tax assistance	366	.377	.868	402	.560	.670	768	.473	.723
Administration	4,169	.658	.763	14,480	.915	.871	18,649	.857	.682
Administration management	1,470	.468	.836	400	.678	.864	1,870	.513	.800
Pay office	100	.790	.746	1,734	.948	.961	1,834	.939	.837
Bookkeeping	328	.811	.869	2,832	.951	.942	3,160	.937	.832
Accounting	975	.461	.915	1,595	.669	.830	2,570	.590	.804
Secretarial work	1,264	.977	.932	5,049	.992	.956	6,313	.989	.893
Office services	18	.778	1.062	1,746	.861	.953	1,764	.861	.953
Clerical work, small firms	14	.857	.758	1,124	.943	.887	1,138	.942	.878
Human resources	1,650	.524	.798	2,953	.858	.813	4,603	.739	.674
HR management	388	.479	.874	95	.663	.968	483	.516	.865
Competence development	417	.511	.887	179	.508	.845	596	.510	.877
Recruiting and employing	279	.616	.820	131	.740	.835	410	.656	.799
Payroll administration	101	.832	.866	1,670	.985	.963	1,771	.976	.800
Safety and health care	336	.351	.710	465	.619	.889	801	.507	.705
Personnel services	129	.713	.847	413	.850	.859	542	.817	.802
Other groups together				9,788	.299	.957	9,788	.299	.957
All	55,499	.221	.877	93,445	.459	.839	148,944	.370	.768

Notes: *N* is the number of observations. Fem is the fraction of female employees in the group. Gap is the sex wage ratio as obtained by dividing the women's mean wage by men's mean wage. 2-digit occupational groups are split into managerial and non-managerial occupations (technical and clerical groups are combined).

We emphasize that the allocation of workers to different jobs is based on a more detailed 3-digit occupation code, which corresponds to the finest classification level. In other words, each occupation group in Table 2.2 includes 1–6 more detailed occupations, resulting in 78 occupations. For example, Office services in Administration include Receptionists, Switchboard Operators, Copyists and Mail Dispatchers, and Office Messengers; and Purchasing in Logistics includes Purchasing Managers, Purchasers, and Purchasing Assistants. Recall that employees of a given firm with the same detailed occupation are assumed to be doing the same job only if they are covered by the same collective agreement, their jobs are located in the same municipality, and in the case of technical and clerical employees their job tasks have been ranked to be of equal worth. Given this narrow definition of the job, we expect that most jobs include only a few employees. We have 5,002 jobs with five or more workers that employ 95,531 workers, and 24,339 jobs with a single worker (see Table 2.3). The number of integrated jobs, i.e. jobs with both sexes present, is 5,018 and 78,631 persons are working in such jobs. We shall discuss the significance of the single worker and fully segregated jobs in the section where we consider how sensitive our results are with respect to the inclusion or exclusion of these groups.

Table 2.3 Distribution of jobs across different size categories

# of women in job	# of men in job							Total	Job size	# of jobs
	0	1	2	3–4	5–9	10–50	> 50			
0	–	12,272	3,126	2,099	1,190	524	19	19,230		
1	12,067	1,128	436	327	321	198	12	14,534	1	24,339
2	2,336	354	171	153	130	108	5	3,257	2	6,590
3–4	1,295	232	93	112	110	121	16	1,979	3–4	4,732
5–9	530	128	75	75	109	119	23	1,059	5–9	2,983
10–50	184	55	45	34	61	118	47	544	10–50	1,799
> 50	3	4	0	3	8	11	31	60	> 50	220
Total	16,415	14,173	3,946	2,848	1,929	1,199	153	40,663	Total	40,663

Notes: 27,670 women and 50,965 men are working in 5,018 integrated jobs; 42,825 men are working in 19,230 jobs with no women; and 27,488 women are working in 16,415 jobs with no men. The total number of jobs is 40,663.

It is obvious that the results of segregation studies are sensitive to the characteristics of the data used. Furthermore, the extent to which our findings are comparable with the results of other studies depends on how similar the analysed worker populations are. We have collected some key characteristics of the data sets used in some other similar, mostly Scandinavian, studies in Table 2.4. Where possible, i.e. in the cases of Norwegian and Danish data, statistics reported refer to white-collar workers in the manufacturing sector. Compared with the other studies, we have the smallest number of occupational categories. However, when we move to the job level by looking at occupations within employers, the discrepancies disappear in the sense that the average job size, i.e. the average

number of workers with the same occupation in the same plant, is quite similar across the data sets. In contrast to the Finnish data, the employer unit in the other data sets is plant, not firm. The largest plant size in the US data results from the constraints required to match a sample of workers to plants. For example, Bayard *et al.* (2003) kept only plants with 25 or more employees in their data. Women's mean wage as a share of men's mean wage ranges from .687 in the US to .768 in Finland, being very similar in Norway, Sweden, and Finland. The share of women is between 30 and 40% in the Nordic data sets. The US data include also blue-collar workers as well as a different set of industries, which may explain the clearly higher female share for the US (47%). In sum, compared with the other Nordic data, our data set seems rather similar in terms of the average job size, sex composition, and gender wage gap.

Table 2.4 Data sets used in similar studies

	This Study	Petersen et al. (1997)	Datta Gupta & Rothstein (2001)	Meyersson et al. (2001)	Bayard et al. (2003)
Country	Finland	Norway	Denmark	Sweden	US
Sector	Manufacturing	Manufacturing	Manufacturing	Private sector	Private sector
Worker group	White-collar	White-collar	White-collar	White-collar	All workers
Sample year	2000	1990	1995	1990	1990
# of workers	148,944	99,486	86,242	391,997	637,718
# of occupations	78	210	98	280	13 / 72 / 491
# of plants/firms	1,464	2,599	2,485	22,031	32,931
# of jobs	40,664	31,692	19,722	146,940	not reported
Mean plant/firm size	102	38	35	18	180
Mean job size	3.66	3.14	4.37	2.67	not reported
Female share	.370	.312	.390	.346	.468
Gender gap	.768	.732	.698	.730	.687

Notes: Only the Finnish data include firms, the other data include plants. The mean plant size for the US data is taken from the last row of column 4 of Table 1 in Bayard *et al.* (2003). The other US figures are sample statistics from the New Worker-Establishment Characteristics Database, where 52 per cent of workers are from manufacturing.

Unfortunately, the other studies lack descriptive characteristics for white-collar manufacturing sector workers. We collected some information on educational attainments of all manufacturing workers in 2000 for Denmark, Sweden, Finland, and the US.²⁰ We found that over 40% of manufacturing workers have received a secondary education in each country. In general, women and men seem to be

²⁰ Information for Swedish and Danish workers was obtained from Statistics Sweden and Statistics Denmark respectively. The US figures were calculated from the 5 per cent Public Use Microdata Sample of U.S. Census Bureau 2000 Census of Population and Housing. The Finnish statistics are from the TT data, i.e. the data used in this study combined with the TT sample of blue-collar workers. Educational information separately for blue- and white-collar manufacturing workers was not available.

allocated rather equivalently across various levels of education, though there may be substantial gender differences in the field of studies.

2.4 Results

2.4.1 Wage regressions

Table 2.5 displays results from the various wage regressions. The explanatory variables are grouped into individual regressors (s_{jki} and \mathbf{x}_{jki}), job regressors ($\bar{s}_{jk} - \bar{s}_{j\cdot}$, $\bar{\mathbf{x}}_{jk} - \bar{\mathbf{x}}_{j\cdot}$, and $\mathbf{g}_{jk} - \bar{\mathbf{g}}_j$), and firm regressors ($\bar{s}_{j\cdot}$, $\bar{\mathbf{x}}_{j\cdot}$, and \mathbf{z}_j). We begin with the fixed effects (FE) model where the full set of job fixed effects (i.e. job dummies) is present. The coefficient of the female dummy takes a value of $-.0627$, indicating that women and men are not equally rewarded by employers. A woman can expect to be paid some 6% lower wages than her equally qualified male co-worker doing the same job within the same firm. One additional year of schooling is estimated to increase the expected wage in a given job by 2.5%. This is clearly below the conventional estimates for the returns to schooling. The much lower estimate obtained here by conditioning on the job held suggests that the wage effect of education works largely through the differential allocation to jobs. Wages increase with tenure and the effect of work experience takes the familiar quadratic form.

The next two specifications include only the fractions of female employees in the set of job and firm regressors. These specifications resemble the specifications estimated by Groshen (1991), Datta Gupta and Rothstein (2001), and Bayard *et al.* (2003). The random effects (RE) specification is estimated with GLS, assuming zero correlation between the latent group effects and variables in \mathbf{x} . As expected, the coefficients of both fraction female variables are negative and statistically highly significant, indicating lower wages for predominantly female firms and jobs mainly occupied by women within firms. It is illustrative to consider the OLS estimates of the same model. The usual OLS standard errors adjusted for heteroskedasticity but derived under the assumption of random sampling are given in the parentheses. In addition, the standard errors that are robust to heteroskedasticity and arbitrary intra-firm correlation are shown in the square brackets (see Wooldridge, 2002, pp. 328–331). The difference between the standard errors is substantial for all coefficients. The usual OLS standard errors that do not take into account the grouped structure of the underlying data are dramatically understated.

Table 2.5 Wage regression results

	Model Specification				
	FE	OLS	RE	CRE1	CRE2
Intercept		1.3483 (.0059) [.0834]	1.8744 (.0083)	1.6734 (.0555)	1.6540 (.0561)
Individual regressors					
Female	-.0627 (.0012)	-.0732 (.0023) [.0054]	-.0653 (.0015)	-.0627 (.0012)	-.0627 (.0012)
Schooling years	.0257 (.0003)	.0886 (.0003) [.0051]	.0425 (.0003)	.0258 (.0003)	.0257 (.0003)
Experience	.0151 (.0002)	.0182 (.0003) [.0013]	.0175 (.0002)	.0151 (.0002)	.0151 (.0002)
Experience ² / 100	-.0228 (.0004)	-.0224 (.0006) [.0040]	-.0254 (.0004)	-.0228 (.0004)	-.0228 (.0004)
$\sqrt{\text{Firm tenure}}$.0138 (.0004)	-.0065 (.0006) [.0038]	.0085 (.0005)	.0138 (.0004)	.0138 (.0004)
Job regressors					
Fraction female		-.1794 (.0028) [.0124]	-.2155 (.0026)	-.1266 (.0022)	-.0779 (.0023)
Mean schooling				.0324 (.0006)	.0218 (.0006)
Mean experience				.0075 (.0005)	.0046 (.0005)
Mean (experience ² / 100)				-.0084 (.0010)	-.0044 (.0010)
Technical				.0317 (.0030)	.0363 (.0029)
Managerial				.3070 (.0022)	.3112 (.0030)
Mean $\sqrt{\text{Firm tenure}}$				-.0102 (.0010)	-.0110 (.0009)
Complexity level					.0457 (.0006)
Complexity missing					.0348 (.0027)
Large city				.0535 (.0029)	.0460 (.0028)
Log (job size)				-.0151 (.0009)	-.0137 (.0008)
Firm regressors					
Fraction female		-.1796 (.0048) [.0572]	-.1407 (.0157)	-.0758 (.0129)	-.0637 (.0130)
Mean schooling				.0280 (.0037)	.0258 (.0038)
Mean experience				-.0006 (.0030)	-.0067 (.0030)
Mean (experience ² / 100)				.0179 (.0064)	.0198 (.0064)
Fraction technical jobs				.0496 (.0127)	.0554 (.0128)
Fraction managerial jobs				.2021 (.0120)	.1667 (.0138)
Mean $\sqrt{\text{Firm tenure}}$				-.0329 (.0030)	-.0321 (.0030)
Mean job complexity					.0099 (.0015)
Fraction complexity missing					.0681 (.0108)
Fraction jobs in large cities				.0441 (.0044)	.0438 (.0044)
Worker mix				.0734 (.0106)	.0688 (.0107)
Mean log (job size)				-.0187 (.0046)	-.0198 (.0048)
Log (firm size)				.0197 (.0028)	.0191 (.0029)
Variance components					
σ_e^2 (individual error)			.0270	.0173	.0173
σ_ω^2 (job random effect)			.0192	.0117	.0101
σ_ξ^2 (firm random effect)			.0115	.0029	.0031

Notes: RE, CRE1, and CRE2 are the (correlated) random effects models with the GLS standard errors in parentheses. The standard errors in parentheses for the fixed effects (FE) and OLS models are adjusted for heteroskedasticity. The FE model corresponds to the model with the full set of job dummies. The standard errors in square brackets for the OLS model are robust to arbitrary heteroskedasticity and intra-firm correlation. All job regressors are measured in deviation from the firm mean. The clerical jobs also include non-managerial jobs in industries where no distinction between the clerical and technical jobs has been made. Models CRE1 and CRE2 include 38 industry dummies. The worker mix variable is the ratio of white-collar employees to all employees. Number of observations is 148,944 in all regressions.

Compared with the fixed effects model, the coefficients of the RE and OLS models are quite different for some individual regressors. Within the random effects framework, this calls into question whether the error components associated with firms and jobs are uncorrelated with the individual explanatory variables. In the presence of such a correlation the regression coefficients in the RE model (and in the OLS model) are inconsistent. In the CRE1 model a number of job and firm controls, including the group means of the individual variables, are added to the model. Testing statistical significance of the coefficients of the group means can be interpreted as a version of the Hausman test for the random effects specification (Wooldridge, 2002, pp. 288–290). Since the coefficients of the group means of x in the CRE1 model are highly significant, we conclude that the RE model is not valid (i.e. the firm and job error components are correlated with x). By implication, the OLS model must be miss-specified as well. In the CRE1 and CRE2 specifications this problem is solved by adding the group means of x to the model, as proposed by Mundlak (1978). Consequently, the coefficients of all individual regressors are identical to those of the FE model.

In the CRE1 model the coefficients of the fraction female variables are reduced by half in absolute value when compared with the RE model. This implies that the fraction female variables in the RE model serve in part as a proxy for other factors responsible for wage differentials between firms and jobs within firms. Coefficients of job and firm regressors generally have signs that seem intuitively reasonable. Larger firms pay higher wages. Firms and jobs within firms that require higher education are associated with higher wages. Workers whose jobs are located in large cities receive better wages, perhaps to compensate for higher living costs. Within firms wage differentials between technical and clerical jobs are relatively small, being around 3% in favour of technical jobs. Managerial jobs are associated with clearly higher wages than technical and clerical jobs.

In the CRE2 specification, the job complexity level is added to the model. This is our preferred model specification. Information on job complexity is missing for all managerial jobs and for some non-managerial jobs. In these cases the dummy for missing job complexity takes a value of one. One additional level of job complexity is found to be associated with a wage increase of almost 5%. Compared with the CRE1 model, the coefficients of managerial and technical jobs remain almost unchanged. If more complex jobs require higher education, it is not surprising to find a considerable fall in the effect of the mean schooling years in job. The coefficient of the fraction female in job is reduced by over one-third in absolute value. Even after controlling for job complexity, average education, and many other factors, it is quite remarkable that wages remain negatively associated with the fraction female variables. In other words, firms and jobs within firms mainly occupied by women pay lower wages for reasons that cannot be explained by observables. This may be due to gender differences in preferences but it may also reflect discrimination through differential access to higher-paying jobs at the point of hire or subsequent promotions. These findings

are partly driven by missing job complexity information for the managerial jobs, however. We address this issue below.

Let us consider the estimates of variance components shown on the bottom lines of the Table 2.5. The variances imply that most of the unexplained wage variation is related to individuals within jobs. In the CRE2 specification the individual error variance σ_ε^2 accounts for about 60% of the total error variance, $\sigma_\varepsilon^2 + \sigma_\omega^2 + \sigma_\xi^2$. The variance of the firm random effects is very low, so that adding more firm-level regressors to the model cannot notably improve the model's fit. The residual intra-job correlation, $(\sigma_\omega^2 + \sigma_\xi^2) / (\sigma_\varepsilon^2 + \sigma_\omega^2 + \sigma_\xi^2)$, describes correlation between the wage outcomes of two randomly chosen workers in a randomly chosen job within a randomly chosen firm, after controlling for s and \mathbf{x} . In the CRE2 model this correlation is as high as .43, which highlights the importance of accounting for the grouped data structure in the econometric modelling.

2.4.2 Wage gap decompositions

Table 2.6 shows the gender wage gap decompositions corresponding to the different model specifications in Table 2.5. From the fixed effect specification we conclude that roughly one-third of the overall gender wage gap of .2505 can be attributed to unexplained within-job wage differentials between sexes (.0627) and gender differences in education, work experience, and firm tenure (.0154). Gender segregation among firms explains some 15% (.0391), while over one-half of the overall gap is owing to gender segregation among jobs within firms (.1334). The OLS and RE models, which were found to be inconsistent, give a different picture about the relative importance of the various determinants of gender wage differentials. The CRE1 and CRE2 models, however, produce decompositions consistent with the fixed effect specification. Contrary to the fixed effects approach, these models allow us to address the question as to why predominantly female firms and jobs are lower paid.

Consider the decompositions associated with the CRE1 and CRE2 models. There are no dominating factors that could be argued to be responsible for most of the aggregate effect of gender segregation among firms (which in turn is quite moderate). Especially, the fraction female in firm has only a minor impact on the gender wage gap, accounting for less than .0080 in both specifications. Likewise, the industry dummies and hence gender segregation among industries do not play any role in explaining the gender wage gap. Gender segregation among jobs within firms is a more interesting case. Over .0500 of the overall wage gap is owing to the disproportionate concentration of men in high-paid managerial jobs. This accounts for about 40% of the aggregate effect of gender segregation among jobs. Among technical and clerical jobs women are more likely to hold less complex jobs, which explains .0256 of the overall gap in the CRE2 decomposition. Once we control for job complexity, the contribution of the

fraction female in job falls from .0688 to .0423, which still accounts for 17% of the overall wage gap.

Table 2.6 Gender wage gap decompositions

	Sample means			Contribution to the wage gap				
	Men	Women	Difference	FE	OLS	RE	CRE1	CRE2
Individual regressors								
Female	.0000	1.0000	-1.0000	.0627	.0732	.0653	.0627	.0627
Schooling years	12.8100	12.0654	.7446	.0192	.0660	.0316	.0192	.0192
Experience	21.3181	22.0808	-.7627	-.0115	-.0139	-.0133	-.0115	-.0115
Experience ² / 100	5.6217	6.0056	-.3839	.0088	.0086	.0097	.0088	.0088
$\sqrt{\text{Firm tenure}}$	3.0579	3.1327	-.0748	-.0010	.0005	-.0006	-.0010	-.0010
				.0781	.1344	.0933	.0781	.0781
Job regressors								
Fraction female	-.2011	.3420	-.5431	.0974	.1170	.0688	.0423	
Mean schooling	.1784	-.3033	.4817			.0156	.0105	
Mean experience	-.0465	.0791	-.1257			-.0009	-.0006	
Mean (experience ² / 100)	-.0242	.0412	-.0654			.0005	.0003	
Technical job	.0334	-.0568	.0901			.0029	.0033	
Managerial job	.0620	-.1054	.1674			.0514	.0521	
Mean $\sqrt{\text{Firm tenure}}$	-.0066	.0112	-.0178			.0002	.0002	
Complexity level	.2075	-.3527	.5602				.0256	
Complexity missing	.0340	-.0579	.0919				.0032	
Large city	.0008	-.0014	.0022			.0001	.0001	
Log (job size)	.1778	-.3023	.4801			-.0073	-.0066	
				.1334	.0974	.1170	.1312	.1304
Firm regressors								
Fraction female	.3318	.4359	-.1042	.0187	.0147	.0079	.0066	
Mean schooling	12.6462	12.3440	.3022			.0085	.0078	
Mean experience	21.3802	21.9753	-.5950			.0033	.0040	
Mean (experience ² / 100)	5.6635	5.9345	-.2711			-.0049	-.0054	
Fraction technical jobs	.1353	.1122	.0231			.0011	.0013	
Fraction managerial jobs	.3992	.3274	.0719			.0145	.0120	
Mean $\sqrt{\text{Firm tenure}}$	3.0445	3.1554	-.1108			.0036	.0036	
Mean job complexity	4.2717	4.0381	.2336				.0023	
Fraction complexity missing	.5345	.4952	.0393				.0027	
Fraction jobs in large cities	.5859	.5698	.0161			.0007	.0007	
Worker mix	.6724	.6855	-.0131			-.0010	-.0009	
Mean log (job size)	2.8923	2.8272	.0651			-.0012	-.0013	
Log (firm size)	6.8822	6.8259	.0563			.0011	.0011	
Industry dummies						.0043	.0051	
				.0391	.0187	.0147	.0380	.0395
Overall sum				.2505	.2505	.2243	.2473	.2480

Notes: The raw wage gap, as measured by the sex difference in mean log wages, is .2505. The first two columns report the sample means of all regressors among men and women; the third column gives the difference. The last five columns show the absolute contribution of each regressor obtained from the various model specifications. The contributions are obtained by multiplying the coefficients in Table 2.5 by the gender differences in sample means in Table 2.1. The cumulative effect of each group of regressors is shown below the horizontal lines.

2.4.3 Robustness of the results

We have carried out several exercises to test the robustness of our results with respect to the model specification and data restrictions adopted.²¹

Missing job complexity information

Information on job complexity is missing for many observations, including all workers in managerial jobs and some workers in technical and clerical jobs. This has two implications. Firstly, the degree of detail in the classification of managerial jobs is low compared with the classification of technical and clerical jobs. This is likely to have an effect on the relative importance of job segregation and that of within-job wage differentials. Secondly, job complexity is only partly controlled for in the CRE2 specification. Therefore, we have performed the same analysis for a sample that excludes observations with missing values for the level of job complexity. This sub-sample represents only technical and clerical workers (some 72,000 individuals), among whom the raw wage gap is much lower (.1570). The magnitude of coefficients in the CRE2 model is different but signs remain unchanged compared with the results in Table 2.5.²² In particular, the coefficient of the female dummy drops to a value of $-.0283$, that of the fraction female in job drops to a value of $-.0335$, and that of the fraction female in firm is reduced by one-third. Compared with the results in Table 2.6, the overall relative contribution of individual-level regressors is lower for the sub-sample. In particular, the unexplained within-job wage gap of .0283 accounts for 18 per cent of the overall gap among technical and clerical workers (compared with 25 per cent for the entire data). The importance of gender segregation among firms appears to be slightly weaker, whereas the contribution of segregation among jobs dominates, accounting for two-thirds of the overall gap. A major part of the effect of gender segregation among jobs within firms (74%) is attributable to gender differences in complexity levels in the jobs held by women and men. This, in fact, explains approximately one-half of the overall gender wage gap among technical and clerical employees.

Varying size thresholds for jobs

Obviously, our definition for the job is quite strict. This leads to a large number of jobs with only one or two workers in the data. One might wonder whether this feature of the data would be partly driving the results. To explore this possibility, we replicated our analysis by excluding all workers in jobs with less than three workers (25% of all observations). This restriction has only a minor effect on the size of the raw wage gap. While the regression results of the CRE2 model remain

²¹ These results are available from the authors on request.

²² See Table 8 in Korkeamäki and Kyyrä (2003).

qualitatively unchanged, the variance of the job random effects shrinks (by 20%) and the coefficients of the job regressors change to some extent. Importantly, changes in the wage gap decomposition turned out to be very small.

Jobs with both sexes present

Another potential problem with our data setup is that apart from being small, quite many jobs consist of only men or women. It is obvious that observations in such jobs contribute to the role of gender segregation, but not to the role of gender wage differentials within jobs. Therefore, we replicated our analysis using a sample restricted to integrated jobs with at least three workers. This restriction results in considerably smaller data. The number of jobs drops from 40,663 to 3,890, the number of workers drops from 148,944 to 76,375, and the raw wage gap shrinks from .2505 to .2070. Considering the CRE2 specification, the coefficients of job regressors are affected but those of individual and firm regressors remain almost unchanged. The most notable difference is more negative coefficients for the fraction female variables. In the wage gap decompositions these changes are reflected in the relative importance of the various sources of gender wage differentials. In the CRE2 decomposition the absolute contribution of job segregation drops from .13 to .05, that of firm segregation rises from .04 to .09, whereas the contribution of unexplained within-job wage differentials and gender differences in individual characteristics remains around .08. The decrease in the importance of job segregation stems from a smaller degree of gender segregation among jobs and from smaller complexity differences between female and male jobs in this sub-sample.

Gender-specific slopes

In the Oaxaca-type decompositions, the regression coefficients are usually allowed to be gender-specific. Table 2.7 shows results from an extended version of the CRE2 model where all individual-level regressors are interacted with the female dummy. For ease of interpretation the education, experience, and tenure variables are measured in deviation from their values for the average worker. The intercept of the model and the female dummy coefficient are shown in the first row of columns 4 and 5 respectively. The other coefficients in column 4 measure returns to education, experience, and tenure for men. Women's returns are obtained as the sum of the coefficients in columns 4 and 5. The female dummy coefficient gives the size of the unexplained within-job gender gap in pay for workers with the average years of education, experience, and tenure. We do not show results for job and firm-level regressors, as they remain unchanged.

Table 2.7 *Regression coefficients and gender wage gap decompositions from the CRE2 model with female interactions in individual-level variables*

	Sample means			CRE2 coefficients		Contribution to the wage gap	
	Men	Women	Difference	Male coefficient	Female interaction	Means	Coeff.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Intercept	1.0000	1.0000	.0000	2.1946 (.0515)	-.0795 (.0018)	.0000	.0795
Schooling years	.2757	-.4688	.7446	.0289 (.0003)	-.0105 (.0004)	.0215	-.0049
Experience	-.2824	.4802	-.7627	.0180 (.0002)	-.0081 (.0003)	-.0137	.0039
Experience ² / 100	.9559	1.3398	-.3839	-.0276 (.0005)	.0127 (.0007)	.0106	-.0171
$\sqrt{\text{Firm tenure}}$	-1.4146	-1.3398	-.0748	.0127 (.0005)	.0026 (.0007)	-.0010	.0035
						.0174	.0649

Notes: From education, experience, and tenure variables, we have subtracted their values for the average worker. Coefficients of the regressors and their interactions with the female dummy are given in columns 4 and 5 respectively. The interaction for the intercept corresponds to the female dummy. The GLS standard errors are in parentheses. The contributions of gender differences in background characteristics in column 6 are obtained by multiplying the coefficients in column 4 with the gender differences in the variable means in column 3. The contributions of gender differences in coefficients in column 7 are obtained by multiplying the interaction coefficients with the sample means for women in column 2 and changing the sign. Cumulative contributions are shown below the horizontal lines.

When β is allowed to differ between sexes, we find lower returns to education and work experience but a higher return to firm tenure for women. In the presence of substantial gender differences in the field of received education, women's lower return to education may indicate that the employers value most technical education, which is received by the majority of men. The absolute value of the female dummy coefficient is slightly higher than in the corresponding specification where β was restricted to be equal for both sexes. Among workers with the average years of education, experience, and tenure, women are found to receive some 7.6% lower wages than their male co-workers who are doing the same job. The expected wage rate increases over the first 35 years of experience but at a lower rate for women. Consequently, the unexplained wage gap within jobs is small for workers with little experience but grows with experience.²³ There is no evidence of the within-job gender gap in favour of men among workers with no experience and tenure; in fact, women may be even slightly better paid in such cases.

²³ Strictly speaking, the within-job wage gap between women and men begins to shrink after 35 years of experience. This phenomenon is driven by the imposed quadratic form for the effect of experience, and does not occur when the experience dummies are used.

The potential years of work experience has a tendency to overestimate actual experience, and the problem is more acute for women because of their higher propensity to be out of work owing to family responsibilities. The imposed functional form for the effect of work experience is also rather rigid. These remarks suggest that our findings may be sensitive with respect to how work experience is measured and incorporated in the model. If we adjust our measure of experience for likely career interruptions and replace the quadratic experience terms with the set of experience terms,²⁴ our main findings remain unchanged, though the female interaction with tenure becomes insignificant.

The last two columns of the Table 2.7 show the contributions of gender differences in the background characteristics (column 6) and regressor coefficients (column 7), i.e. the explained and unexplained part of the wage differential between the average woman and average man who are doing the same work for the same employer respectively. Some 80% of the within-job wage differential of .0823 remains unexplained, being attributed to different coefficients for women and men. Only a small fraction of the gender wage differential is attributable to gender differences in the background characteristics, namely to women's lower education level. Not surprisingly, these are the same conclusions we jumped to in the case of the model without gender-specific slopes.

It is important to recognize that the results of decomposition exercises with gender-specific slopes are somewhat arbitrary. Firstly, the male coefficients are taken as a reference structure that is used to evaluate the contributions of gender differences in the background characteristics in Table 2.7. Alternatively, one could choose the female coefficients or some weighted average of male and female coefficients. If we use the female coefficients as the reference structure, our main insights do not change, though the relative contribution of women's lower education level is reduced. Secondly, the estimated contributions of gender differences in the slopes are affected by the location transformations of the explanatory variables (Oaxaca and Ransom, 1999). We found that the unexplained within-job wage differential in the last column is attributed mainly to the female dummy coefficient, while women's lower returns to education and experience have only moderate contributions. These findings are driven by our choice to measure the education, experience, and tenure variables in deviation from their values for the average worker. If these variables were transformed to range from zero upwards, we would find a negligible (or slightly negative)

²⁴ From men's experience we subtracted one year for military service. We subtracted 1, 2, or 3 years from women's experience if the original experience variable was 2–4, 5–7, or more than 7 years respectively. These modifications are based on Asplund's (2001) comparisons of the potential and actual (self-reported) years of work experience.

female dummy effect and strong positive contributions for women's lower returns to education and experience. In such a case, the female dummy coefficient would pick up the size of the within-job gender gap in pay among the lowest educated workers with no experience and tenure. In other words, the results in the last column depend on the interpretation we put on the female dummy coefficient. Importantly, the overall effect of gender differences in the regression coefficients is not affected by the location transformations, or the estimated contributions of gender differences in the background characteristics in column 6.

Despite these difficulties in interpretation, it is safe to make the following conclusions: 1) only a small fraction of the within-job wage gap between sexes is attributable to women's lower education level, while most of the gap remains unexplained. 2) The unexplained within-job wage differential increases with education and experience, being very small or negligible among low educated workers with little experience and tenure.

2.4.4 Comparisons with findings from other studies

We conclude this section by contrasting our main findings with the findings from other studies that explore the importance of labour market segregation and wage differentials occurring within jobs (i.e. within occupations within plants) in explaining the gender wage gap.²⁵ There are only a few such studies, as the evaluation of within-job wage differentials calls for high quality matched employer-employee data, which are not widely available. Our results are most directly comparable with evidence for white-collar workers in Norway (Petersen *et al.*, 1997), Sweden (Meyersson Milgrom *et al.*, 2001), and Denmark (Datta Gupta and Rothstein, 2001).²⁶ We found that in Finland white-collar women earn some 22% less on average than men do, compared with 30% in Denmark and 27% in Norway and Sweden. Gender segregation among firms accounts for about 16% of the Finnish raw gap, while the industry effects have no significant role at all. These results are consistent with the findings for other Nordic countries where gender segregation among industries or employers does not play an important role in the case of white-collar workers. Roughly one-half of the

²⁵ The most up-to-date and comprehensive study of gender wage differentials in Finland that does not go to the job-level is Vartiainen (2002). It includes the standard Oaxaca decompositions for workers in the private service sector, for white- and blue-collar manufacturing workers, and for local and central government workers.

²⁶ The Norwegian data of Petersen *et al.* (1997) contain (almost) all workers in six business sectors in 1984 and 1990. The data used by Meyersson Milgrom *et al.* (2001) is more extensive, covering most privately employed workers in Sweden over the period 1970–1990. Datta Gupta and Rothstein (2001)'s Danish data include (almost) all salaried workers in 1984 and 1995 in the private sector. In the text we refer only to the results for the latest period available, and for Norway and Denmark we report figures for white-collar manufacturing workers.

raw gap of the Finnish white-collar workers is attributable to the disproportionate concentration of women in lower-paying jobs within firms. In Denmark occupational and job segregation explains roughly one-half of the raw gap but in Norway and Sweden a much higher share. Finally, we found that within jobs white-collar women are paid some 6% lower wages on average than their male co-workers with equal education, work experience, and tenure. This figure is roughly identical to the size of the (unconditional) within-job wage gap found for Sweden and Norway but much lower than what has been found for Denmark (about 14% after controlling for a number of individual characteristics).

In Korkeamäki and Kyrrä (2003) we also analyse blue-collar workers in the Finnish manufacturing sector, among whom women's mean wage is 16% lower than men's mean wage. Petersen *et al.* (1997) and Meyersson Milgrom *et al.* (2001) report somewhat lower wage gaps for blue-collar workers in Norway and Sweden respectively. We found that most of the gender gap among blue-collar workers results from gender segregation among firms. This is in accordance with a strong effect of establishment segregation among the Swedish blue-collar workers.²⁷ By contrast, Petersen *et al.* (1997) find employer segregation less important in the case of Norway. Furthermore, we found that blue-collar women are paid 3.5% less than their equally qualified male counterparts who are doing the same job for the same employer. This figure is very close to the (unconditional) within-job gap in Norway, being above the Swedish level.

In a complementary study, Luukkonen (2003) explores gender wage differentials in the Finnish service sector using the methods developed in this paper. White- and blue-collar workers were not analysed separately. Half of the overall gap of 20% in the service industries results from gender segregation among jobs, one-third from gender segregation among firms, and one-sixth is owing to unexplained within-job wage differentials and gender differences in individual characteristics. Within-job wage differentials between sexes remain mostly unexplained, women being paid 3.7% less for the same job than their equally qualified male co-workers.

Comparisons with US evidence are less straightforward because the US findings are mixed and because the US studies do not make a clear difference between white-collar and blue-collar workers. Groshen (1991) and Petersen and Morgan (1995) find that the segregation of women into lower-paying occupations, industries, and establishments essentially explains all of the gender wage

²⁷ Carrington and Troske (1998) find that in the US manufacturing sector gender segregation among plants is more extensive and explains a higher fraction of the gender wage gap among blue-collar workers than among white-collar workers.

differentials in the US labour market.²⁸ By contrast, Bayard *et al.* (2003) report large and significant within-job wage differentials in the US labour market, even though gender segregation accounts for at least one-half of the gender wage gap.²⁹ Compared with Bayard *et al.* (2003), our results point to a smaller (unexplained) within-job gender gap and a stronger role for gender segregation in Finland. These conclusions are reversed, if the findings of Groshen (1991) or Petersen and Morgan (1995) are taken as a reference.

2.5 Conclusion

In this paper, we have introduced a new way of decomposing the gender wage gap based upon the correlated random effects model. The decomposition allowed us to assess the extent to which the overall gender gap is attributable to within-job wage differentials and gender segregation among firms and jobs within firms. Importantly, by explicitly modelling the firm and job effects, the approach proved to be informative about the sources of lower pay in predominantly female firms and jobs. Compared with the standard fraction female decomposition, the correlated random effects specification led to quantitatively different results. These differences suggest that the latent firm and job effects may bias the coefficients in the simple fraction female regressions and lead to misleading conclusions.

A major part of the gender wage differentials among white-collar manufacturing workers was attributed to the disproportionate concentration of women in lower-paying jobs. Within firms, high-paid managerial jobs are mainly occupied by men, and among other types of jobs men are concentrated in positions with higher skill requirements. This may reflect discrimination through differential access to higher-paying jobs, or it may result from gender differences in preferences. Becker (1985), for example, illustrates how women's greater responsibility for childcare and housework may induce them to crowd into less-demanding jobs, as well as to expend less effort in the same job than men do. Although the reasons for the preponderance of women in lower-paying jobs remain a puzzle, our findings highlight the importance of equal opportunities in education, hiring and promotion.

When explaining wage differentials between jobs, we found that lower wages in predominantly female jobs are in large part attributable to lower skill requirements and job complexity (especially when managerial jobs were

²⁸ The data used by Groshen (1991) and Petersen and Morgan's (1995) cover a narrow subset of occupations in few industries.

²⁹ While not entirely representative, the data used by Bayard *et al.* (2003) include workers and establishments from all sectors of the US economy. In the data, women's mean wage is 31 per cent lower than men's.

excluded). However, our results suggest that, in the same firm, predominantly female jobs pay lower wages than predominantly male jobs that are associated with the same level of average education, average tenure and job complexity. In other words, jobs of equal worth are differently rewarded depending on whether they are occupied by men or women. Of course, one can always speculate how accurate our job complexity variable is, but if we assume that this measure is reasonably good, our results would imply that policies like comparable worth might be worth considering. This does assume that men and women exert equal effort in these jobs (e.g. Becker, 1985). Even if effectively implemented, the scope of such measures is likely to be relatively limited, however, because the share of lower wages in predominantly female jobs that was left unexplained accounts for less than a fifth of the overall gender gap. Finally, we found that within jobs women are paid some 6% less than their equally qualified male co-workers. The unexplained within-job gap is higher among more educated and more experienced workers. Eliminating the sources of unexplained within-job wage differentials can directly account for a quarter of the overall gender gap in pay.

References

- Altonji, J. – Blank, R. (1999): Race and gender in the labor market. In Ashenfelter, O. – Card, D. (eds.): *Handbook of labor economics*, Vol. 3C. Elsevier Science, Amsterdam, 3143–3259.
- Asplund, R. (2001): Education and earnings in Europe: Finland. In Halmon, C. – Walker, I. – Westergård-Nielsen, N. (eds.): *Education and earnings in Europe: a cross-country analysis of the returns to education*. Edward Elgar, Cheltenham, 68–83.
- Bayard, K. – Hellerstein, J. – Neumark, D. – Troske, K. (2003): New evidence on gender segregation and gender differences in wages from matched employee–employer data. *Journal of Labor Economics* 21(4), 887–922.
- Becker, G. S. (1985): Human capital, effort, and the sexual division of labor. *Journal of Labor Economics* 3(1), S33–S58.
- Blau, F. – Kahn, L. (2000): Gender differences in pay. *The Journal of Economic Perspectives* 14(4), 75–100.
- Carrington, W. J. – Troske, K. R. (1998): Gender segregation in U.S. manufacturing. *Industrial and Labor Relations Review*, 51(3), 445–464.
- Chamberlain, G. (1984): Panel data. In Griliches, Z. – Intriligator, M. D. (eds.): *Handbook of econometrics*, Vol. II. Elsevier Science, Amsterdam, 1247–1318.

- Datta Gupta, N. – Rothstein, D. S. (2001): The impact of worker and establishment-level characteristics on male-female wage differentials: Evidence from Danish matched employee-employer data. Centre for Labour Market and Social Research Working Paper 01-09-2001.
- Fuller, W. A. – Battese, G. E. (1973): Transformations for estimation of linear models with nested-error structure. *Journal of the American Statistical Association* 68(343), 626-632.
- Groshen, E. L. (1991): The structure of the female/male wage differentials: Is it who you are, what you do, or where you work? *Journal of Human Resources* 26(3), 457-472.
- Korkeamäki, O. – Kyyrä, T. (2002): The gender wage gap and gender segregation in Finland. VATT Discussion Papers 288.
- Korkeamäki, O. – Kyyrä, T. (2003): Explaining gender wage differentials: Findings from a random effects model. VATT Discussion Papers 320.
- Luukkonen, A. (2003): The gender wage gap in private service occupations (in Finnish). VATT Discussion Papers 321.
- Lilja, R. (1995): Career mobility in Finnish industry. ETLA Discussion Papers No. 544.
- Meyersson Milgrom, E. – Petersen, T. – Snartland, V. (2001): Equal pay for equal work? Evidence from Sweden and a comparison with Norway and the US. *Scandinavian Journal of Economics* 103(4), 559-583.
- Moulton, B. R. (1986): Random group effects and the precision of regression estimates. *Journal of Econometrics* 32, 385-397.
- Mundlak, Y. (1978): On the pooling of time series and cross-section data. *Econometrica* 46(1), 69-85.
- Oaxaca, R. L. – Ransom, M. R. (1999): Identification in detailed wage decompositions. *The Review of Economics and Statistics* 81(1), 154-157.
- Petersen, T. – Morgan, L. A. (1995): Separate and unequal: Occupation-establishment gender segregation and the gender wage gap. *American Journal of Sociology* 101(2), 329-365.
- Petersen, T. – Snartland, V. – Becke, L.-E. – Modesta Olsen, K. (1997): Within-job wage discrimination and the gender wage gap: the case of Norway. *European Sociological Review* 13(2), 199-213.
- Vartiainen, J. (2002): Gender wage differentials in the Finnish labour market. Labour Institute for Economic Research Discussion Papers 179.
- Wooldridge, J. M. (2002): Econometric analysis of cross section and panel data. The MIT Press, Cambridge.

3. A distributional analysis of earnings losses of displaced workers in an economic depression and recovery³⁰

Abstract

We study the earnings losses of Finnish private sector workers who lost their jobs at two very different points in the business cycle. The first group was displaced in 1992 (depression period) and the second one in 1997 (recovery period). The focal point of the analysis is the quantile displacement effect, the change in the earnings distribution due to involuntary job separation. We use mass layoffs and plant closures to identify groups of workers who were displaced for exogenous reasons. The effect of displacement is strongest at the lower end of the earnings distribution, and small or negligible at the upper end. Women and those displaced during the depression period were subject to the largest earnings losses.

Keywords: Displacement, earnings losses, unemployment, quantile regression.

JEL classification: J31, J63, J65

3.1 Introduction

In all labour markets, a large number of workers lose their jobs every year. Some job losers are re-employed quickly without significant earnings losses. Others remain unemployed for long periods or have to accept large cuts in wages, or they may be pushed out of the labour market. Job displacement can lead to substantial individual costs in terms of foregone earnings and employment. These costs have been the focus of a number of recent studies. Evidence from US studies suggests that the average earnings losses of displaced workers are large and persistent, being around 10–25 per cent even several years after the job loss (see Ruhm, 1991, Jacobson *et al.*, 1993, Stevens, 1997, and Couch and Placzek, 2010). However, the reduction in employment following displacement has been found to be relatively short-lived in the US labour market. Some studies,

³⁰ Ossi Korkeamäki and Tomi Kyyrä.

We are grateful to Jukka Appelqvist for his help with the data and preliminary analysis. We thank Kristiina Huttunen, Matti Sarvimäki, Jouko Verho, and Ralf Wilke, as well as the participants at the EALE Conference in Oslo, the ESPE Conference in Chicago, the EcoMod Conference in Sao Paolo, the Annual Meeting of the Finnish Society for Economic Research in Lappeenranta and the HECER Workshop on labour turnover and firm performance in Helsinki for helpful comments.

including Couch (2001), Burda and Mertens (2001), and Bender *et al.* (2002) for Germany, Huttunen *et al.* (forthcoming) for Norway, and Hijzen *et al.* (2006) for the UK, suggest that the long-term costs of job loss are small or non-existent in the European labour markets. On the other hand, studies by Borland *et al.* (2002) for the UK, Margolis (1999) for France, Carneiro and Portugal (2006) for Portugal, Eliason and Storrie (2006) for Sweden, and Appelqvist (2007) for Finland find the long-term losses to be much larger and more concordant with the findings from the US. Although the results from these studies are not directly comparable due to the different time periods analysed and large dissimilarities in the underlying data and research design, there seem to be significant differences in the displacement cost between countries.

Traditional analysis of earnings losses associated with displacement has typically employed classical least-squares regression methods. Although the resulting effect on the conditional mean of earnings is of considerable interest, the distributional aspects of earnings losses are equally important. Earnings dispersion provides, for example, a measure of uncertainty about future earnings. For a given mean loss, a larger increase in dispersion following displacement implies a larger welfare loss for the risk-averse worker. Furthermore, the mean impact is not indicative of the size and nature of the effect of displacement in the tails of the earnings distribution, which might be of primary interest from the policy point of view. A strong negative effect in the upper tail would suggest lower chances of being re-employed in a highly paid job for the displaced worker. An equally strong effect in the lower tail is perhaps more alarming, because it would imply a high risk of joblessness or low-paid employment. If such an effect still exists several years after the displacement period, it may call for directed supportive measures. When the focus of the analysis is restricted solely to the mean impacts, these pieces of information will remain missing. In general, the change in the conditional mean gives an incomplete picture of the consequences of displacement. A more complete picture can be obtained by estimating a family of conditional quantile functions, which is the approach we take in this study.

This study considers the effect of job displacement over the entire distribution of earnings in Finland. We use linked employer-employee panel data to construct groups of private sector employees who lost their jobs at two very different points in the business cycle. The first group was displaced in 1992 (depression period) and the second group in 1997 (recovery period). Following the standard practice, we take separations associated with mass layoffs and plant closures to be job displacements, as they are likely to be exogenous from the workers' standpoint. These groups of displaced workers and the associated comparison groups are followed over an 11-year period beginning three years before and ending seven years after the year of possible displacement. To include all the costs of job loss we also include periods with no earnings resulting from long-term unemployment or non-participation in the analysis.

Our two follow-up periods reflect markedly different macroeconomic conditions. At the end of the 1990s, the Finnish economy grew at a high rate and unemployment was declining. By contrast, GDP dropped over 10% between 1991 and 1993, causing the unemployment rate to exceed 16%. As a result, workers who were displaced in 1992 lost their jobs in the middle of one of the deepest recessions in advanced countries since the 1930s. A few years ago, one might have viewed the depression that hit the Finnish economy in the early 1990s as an extraordinary event, being comparable only to the experiences of a few other countries during the Great Depression of the 1930s. Nevertheless, because of the current global economic crisis triggered by the US subprime mortgage collapse in 2007, several countries across the world are experiencing production and employment losses akin to those seen in Finland in the early 1990s. By analysing the consequences of displacements in the Finnish depression, we get at least a rough picture of the size and duration of the earnings losses that millions of people losing their jobs during the current crisis are likely to experience in the near future.

According to our results, earnings losses are especially large when job loss occurs in the depression period. In that case, the entire earnings distribution still lies below the counterfactual distribution (without job loss in the reference period) seven years after the job loss. Losing a job in a recovery period also has a long-lasting effect, but only at the lower end of the distribution. Women tend to suffer from larger earnings losses irrespective of the timing of job loss in the business cycle. We also find evidence of considerable heterogeneity in the displacement effect. The effect of job loss is concentrated in the lower end of the distribution, being relatively moderate (the depression period) or even negligible (the recovery period) at the upper end. Finally, we show that job loss not only causes a decline in expected earnings, but also raises uncertainty about the future earnings level due to a substantial increase in the earnings dispersion.

The rest of the paper proceeds as follows. In the next section, we discuss the evaluation issues and define the quantile displacement effect. Section 3.3 describes our data and the selection of various worker groups for the regression analysis. Descriptive evidence is presented in section 3.4, which is followed by the quantile regression results in section 3.5. The final section concludes.

3.2 Evaluation issues

We are interested in the effect of job displacement in a past period on the current earnings. Since the involuntary job loss can be viewed as a “treatment”, we can discuss the evaluation issues along the lines of the extensive literature on programme evaluation. Let Y_1 be the earnings in the current period if the worker was displaced in a reference period, and let Y_0 be the earnings in the counterfactual situation without displacement in the reference period. If we could measure workers earnings in both states, the displacement effect, $Y_1 - Y_0$ would

be directly observed. The fundamental problem of causal inference is that both Y_1 and Y_0 are never observed for the same individual (Holland, 1986).

The observed earnings can be expressed as

$$Y = DY_1 + (1 - D)Y_0, \quad (14)$$

where $D = 1$ if the worker was displaced, and $D = 0$ otherwise. Displacements are certainly not randomly assigned but dismissed workers are selected in a complicated procedure that takes into account individual characteristics, some of which may not be observable for the researcher. This implies dependence between D and (Y_1, Y_0) even after controlling for a wide array of individual characteristics. For example, a good match between a worker and a job in the reference period may imply a high value of Y_0 and a low probability of displacement. One consequence of this is that simple comparisons of outcomes between displaced and non-displaced workers do not have a causal interpretation.

A vast majority of displacement studies have exploited mass layoffs or plant/firm closures to detect workers who lost their jobs from exogenous reasons. Such workers are less likely to be laid off because of their own characteristics or performance, but because of a shock that hit their employer. Under this assumption, one can overcome the endogeneity problem by choosing a sample where displacements result from mass layoffs or plant/firm closures. This approach is taken also in this study. Thus D is hereafter an indicator of a displacement associated with a mass layoff or a plant closure, and it is assumed to be independent of (Y_1, Y_0) given observed individual characteristics.

3.2.1 Expected earnings losses

Since the individual effects of displacement cannot be identified, the focus of the evaluation literature has been on estimating the average effect, $E(Y_1 - Y_0)$, or the average effect on the displaced, $E(Y_1 - Y_0 | D = 1)$. Let us assume that $E(Y_0 | \mathbf{X} = \mathbf{x}) = \mathbf{x}'\beta$ and $Y_1 = Y_0 + \alpha$, where α is an individual-specific effect of displacement which is independent of D and \mathbf{X} . Under these assumptions we could estimate $\delta = E(\alpha) = E(Y_1 - Y_0) = E(Y_1 - Y_0 | D = 1)$ by least squares from

$$y = \mathbf{x}'\beta + \delta d + \varepsilon. \quad (15)$$

The estimates of displacement costs have usually been obtained from regressions similar to this stylized example.³¹ There are some pitfalls worth noting. First, the

³¹ Displacement studies have usually exploited panel data. The estimating equations are somewhat more involved than (2) due to repeated observations on the same workers over time. With panel data, one can

data on earnings typically involve observations of zero earnings for the long-term unemployed and those who withdrew from the labour market. It follows that the outcome variable takes on the value zero with a positive probability (In other words, there is a mass point at zero) but is continuously distributed over strictly positive values. A common practice in the displacement cost literature has been to restrict the analysis to a subset of observations with strictly positive earnings. This results in a selective sample of those who were able to return to work after displacement, and hence the estimated effects of displacement are potentially subject to selection bias. Angrist (2001) argues that, in the context of limited dependent variables, estimates obtained using positive outcomes only do not have a meaningful causal interpretation as treatment effects even if the data come from an ideal randomized experiment. Keeping zero observations in the analysis and applying least squares to the full sample is arguably a better alternative, even though the conditional mean is unlikely to be linear in \mathbf{x} and d . Alternatively, non-linear models for the conditional mean of limited dependent variables, like Tobit or sample selection models, can be used, but such models rely on strong parametric assumptions (about homoskedasticity, symmetry or functional forms) that may be difficult to justify in practice.

The displacement effect in (15) is also assumed to be homogeneous in the sense that it shifts the *location* of the earnings distribution (though possibly in an individual-specific way) without affecting other distributional aspects, such as dispersion, skewness or tail behaviour. The ranking of workers in the earnings distribution (conditional on $\mathbf{X} = \mathbf{x}$) is, however, partly determined by past luck and success in the labour market. In the theoretical models of job search, employed workers are looking for better jobs and climb up the job ladder when a higher-paying job is found. When this time-consuming process is interrupted by involuntary job loss, the worker has to restart the job search from the bottom. Search theory suggests that the upper end of the distribution of Y_0 given $\mathbf{X} = \mathbf{x}$ is disproportionately populated by workers who have been lucky to find good jobs through several switches to better jobs over a long period. In the case of job loss, these workers are likely to experience larger earnings losses than their less lucky counterparts at the lower end of the distribution of Y_0 who would probably be employed in bad jobs also without a displacement. In other words, the effect of displacement increases with Y_0 given $\mathbf{X} = \mathbf{x}$. Here the displacement effect heterogeneity stems from random events that may be independent of workers' characteristics.

Alternatively, the ranking of observationally identical workers in the earnings distribution may reflect individual-specific characteristics not observed by the

also allow the displacement status to be correlated with the error term by introducing fixed individual effects. These panel data models are subject to the same pitfalls discussed in the text.

researcher. The upper end of the distribution of Y_0 (conditional on $\mathbf{X} = \mathbf{x}$) may be populated by high-ability workers who are able to return to work quickly after a job loss at a wage rate close to their previous wage. In contrast, those at the lower end of the earnings distribution without displacement may be less able workers who would have trouble in finding work after displacement, and hence are subject to potentially large earnings losses due to long periods out of work if displaced. This is consistent with the hypothesis that the relative importance of transferable individual-specific skills, which are valuable also in a new job following displacement, is higher for high ability persons who can qualify for high-end occupations.³² This kind of reasoning would suggest that the effect of displacement decreases with Y_0 given $\mathbf{X} = \mathbf{x}$.

In general, the displacement effect is likely to be heterogeneous, in which case the mean impact does not tell the whole story. The quantile regression method of Koenker and Bassett (1978) for distributional analysis provides an alternative approach.

3.2.2 Distributional analysis

To define our displacement effect we follow the quantile treatment literature, going back to Lehmann (1974) and Doksum (1974). Let F_{y_1} and F_{y_0} be the cumulative distribution functions of Y_1 and Y_0 respectively. We define the *quantile displacement effect* (QDE) at the θ -th quantile as

$$\alpha_\theta = F_{y_1}^{-1}(\theta) - F_{y_0}^{-1}(\theta), \quad (16)$$

where $F_{Y_j}^{-1}(\theta) = \inf\{y_j \mid F_{Y_j}(y_j) \geq \theta\}$, $j = 0, 1$, for $\theta \in (0, 1)$. In other words, α_θ equals the horizontal distance between the distribution functions of potential earnings with and without displacement at given θ . A family of α_θ over θ captures heterogeneity in the displacement effect over the distribution of potential earnings. More precisely, what is captured is the difference between the two *marginal* distributions. For example, $\alpha_{0.5}$ describes the difference in the median earnings with and without displacement, not the effect of displacement on the earnings of a worker with median earnings in the absence of displacement. The difference in the distributions is all we can identify from the observed data, without imposing strong additional restrictions. Nevertheless, the QDE estimates can be very informative, as they reveal whether job displacement reduces expected earnings (the distribution shifts left), increases uncertainty about the

³² Kletzer (1989) finds that a larger part of returns to tenure is individual-specific as opposed to match- or job-specific, and hence not lost in job displacement, for managerial, professional and technical workers than for blue-collar workers in the US labor market.

future earnings (dispersion increases), or has different effects at the lower and upper ends of the distribution.

In the absence of covariates, the natural and simple estimator of the QDE is obtained by replacing $F_{Y_1}^{-1}(\theta)$ and $F_{Y_0}^{-1}(\theta)$ with their empirical counterparts. This would require two *randomized* samples of individuals: those who were displaced in the reference period and those who were not. Then, for example, the difference in median earnings in the current period between the displaced and non-displacement groups would give an estimate of the QDE at $\theta = 0.5$. We do not believe that our sample design based on mass layoffs and plant closures is comparable to a randomized experiment, but we do assume independence of the displacement status conditional on the control variables. Since this assumption is crucial for causal interpretation of the estimated displacement effects, we shall provide (indirect) empirical evidence to support its validity in our application.

Koenker and Bassett (1978) introduced the quantile regression method for estimating conditional quantile functions. Powell (1986) developed an estimator for the conditional quantiles of limited dependent variables. We parameterize the conditional quantiles of the potential earnings as:

$$\begin{aligned} F_{Y_0}^{-1}(\theta | \mathbf{x}) &= \max\{0, \mathbf{x}'\beta_\theta\}, \\ F_{Y_1}^{-1}(\theta | \mathbf{x}) &= \max\{0, \mathbf{x}'\beta_\theta + \alpha_\theta\}, \end{aligned} \tag{17}$$

where the limited support of the earnings distributions is explicitly accounted for. Provided that D is independent of (Y_1, Y_0) given $\mathbf{X} = \mathbf{x}$, the conditional quantile function for observed earnings can be written as

$$F_Y^{-1}(\theta | \mathbf{x}, d) = \max\{0, \mathbf{x}'\beta_\theta + \alpha_\theta d\}. \tag{18}$$

This type of model can be applied to corner solution data or to censored data with the fixed censoring point. In our application, the issue is *not* data observability: the earnings are observed for all workers but they are zero for those who did not work in the period in question. In other words, the outcome variable is not censored but has a mass point at zero, being continuously distributed over strictly positive values. Wooldridge (2002) calls models for such outcome variables *corner solution* models. These are statistically identical to models for censored data, but are conceptually very different, which should be kept in mind when interpreting the results. Under the corner-solution interpretation, the conditional θ -th quantile of Y is zero for $\mathbf{x}'\beta_\theta + \alpha_\theta d \leq 0$, while it is strictly positive and linear

in \mathbf{x} and d for $\mathbf{x}'\beta_\theta + \alpha_\theta d > 0$.³³ In the former case, the worker is predicted to be out of work – and hence has zero earnings – with a probability no less than θ .

Since Powell's (1986) estimator does not require additional parametric assumptions, we can recover the effect of displacement over the conditional distribution in a robust way despite the limited support of the outcome variable. In summary, we apply the same identifying assumption – the conditional independence of D – that has been commonly assumed in linear specifications like (15), but we explicitly model heterogeneity in the effect of D (which is of particular interest) without excluding observations with zero earnings from our analysis (which would lead to the selection bias).³⁴

3.3 Data and sample construction

Our data come from the Finnish Longitudinal Employer-Employee Database of Statistics Finland. This database combines information from several administrative registers for all working age persons with a permanent residence in Finland. It includes detailed information on employment and earnings history along with a number of background characteristics, like education, marital status and age. Since the data include all people, not just those who are currently employed, we can follow individuals irrespective of their labour market state, provided they have not emigrated or died. This is important because job loss may be followed by periods of unemployment and non-participation. For example, a worker who loses his or her job in a sunset industry may withdraw from the labour force temporarily in order to acquire new skills required by jobs in other industries. Thus, if all costs of job loss are to be included, we should not exclude periods out of work from our analysis. The database also includes unique identification codes for all plants (and firms) operating in Finland. This information allows us to detect all employees of a given plant at the end of any given year, as well as to identify plants that were downsizing or exiting the market in a given year.

Starting from Ruhm (1991) and Jacobson *et al.* (1993), practically all of the recent studies on the cost of job loss have employed a methodology that involves a comparison of displaced workers with a control group that did not experience

³³ It should be stressed that we view that the limited support of the dependent variable is a technical problem. The mass point in the earnings data at zero implies that the linearity assumption of the conditional quantile function is not valid in the left tail of the distribution, whereas the consequences for higher regression quantiles are negligible. In our application, the standard quantile regression method, which ignores the limited support of the dependent variable, would lead to highly similar estimates to the corner-solution approach except for the first decile, which is the lowest quantile analyzed.

³⁴ Carneiro and Portugal (2006) also apply a quantile regression model to study earnings losses in Portugal but they exclude observations with zero earnings from their analysis.

displacement during a given reference period. Displacements are typically defined as permanent and involuntary separations caused by an employer-specific shock, not related to the worker's job performance. In practice, it is not possible to distinguish directly between layoffs and quits on the one hand, and between employer-specific and individual-specific reasons for separations on the other. A common solution in empirical work has been to interpret separations associated with a mass layoff or a plant/firm closure to be displacements. The underlying assumption is that such separations are driven by employer-specific shocks, and hence exogenous from the worker's standpoint.

Obviously, this strategy is not completely accurate. First, displacements defined in this way may also include some voluntary quits. Second, the employer has an incentive to get rid of the least productive workers first, although the seniority rules and unions may prevent the employer from choosing freely the group to be laid off. This suggests that workers who are displaced in a mass layoff are probably not a completely random group. Therefore, some researchers prefer the use of plant (or firm) closures to mass layoffs. A counter-argument is that those plants that closed down are a more selective group of all plants (for example, they are much smaller on average) than downsizing plants are, suggesting that their employees may also be a rather selective group. In the absence of a superior solution, we include both groups in our analysis, but estimate distinct displacement effects for those who lost their jobs in mass layoffs and for those who were displaced due to plant closures.

Finally, the plant closure or mass layoff is likely to be expected by employees, and thereby some of them may quit earlier in anticipation of the forthcoming reduction in the workforce. If workers with better outside options are more likely to leave early, those who are displaced in the year of the mass layoff or plant closure form a selective group. On the other hand, the downsizing process can be longer than one year, so that the employer may have laid off some workers well before the period of the mass layoff or plant closure. This suggests that some of the early leavers may be low productivity employees. One could classify the early leavers as displaced workers, but then more voluntary quits and more selective dismissals would be included as well. Therefore, we instead include workers who left their jobs a year before a mass layoff or plant closure as a separate group in our analysis.

We construct two separate samples using 1992 and 1997 as base years when the event of displacement possibly took place. We focus on workers who all have initially a strong labour market attachment, and thereby require that everyone included in the sample have at least three years of tenure with the same private sector employer before the base year. We also require that during these three years everyone included in the sample had exactly one employer, were not self-employed, and did not have any unemployment spells. The employers are identified using plant codes and we only include workers from plants that employ

at least ten workers at the end of the year preceding the base year.³⁵ Furthermore, we require that all workers in the sample were 20 to 51 years old in the base year.³⁶

For both base years, we identify a group of displaced workers (the *displacement* group) as well as a group of workers who were not displaced at that time (the *control* group). The control group includes employees who did not separate from their employer during the base year, In other words, had the same plant code at the end of the base year as they had a year earlier. The displacement group consists of two subgroups: those who lost their jobs in mass layoffs and those who were displaced due to plant closures. The former subgroup includes all workers who separated during the base year from plants from which at least 50% but not all of their employees left by the end of the base year.³⁷ The latter subgroup consists of separating workers whose plant disappeared entirely (in terms of employment) by the end of the base year. From the displacement group we exclude workers who return to the pre-displacement plant or firm at some later period (cf. our definition of displacement). We also include a group of workers who, during the base year, left their jobs in plants that downsized or closed down in the *next* year (the *early-leaver* group). These workers are analyzed separately and the results for this group are used as a robustness check. To summarize, we have defined the following groups for the base year t (1992 or 1997):

Control group: Workers who did not change their employer during year t .

Displacement group: Workers separating in year t from plants that closed down during year t (*plant-closure* subgroup) and from plants that reduced workforce at least by 50% between years $t - 1$ and t but were still in operation at the end of year t (*downsizing* subgroup).

Early-leaver group: Workers separating in year t from plants that closed down during year $t + 1$ (*plant-closure* subgroup) and from plants that reduced their

³⁵ We consider plants, not firms, as production units that are subject to a risk of downsizing and closure. In doing so, we avoid problems with artificial firm closures that result from changes in the firm identifiers due to mergers or dispersals and changes in ownership or industry classification. The plant codes do not suffer from the same problems, as the plant is defined as a local kind-of-activity unit in the underlying register data.

³⁶ By excluding individuals over 51 years of age, we rule out the possibility of early retirement via the unemployment tunnel scheme that consists of extended unemployment benefits and a particular unemployment pension scheme for the older long-term unemployed. The older unemployed entitled to this scheme are a very distinctive group, as roughly half of them have been estimated to have effectively withdrawn from job search and to be waiting passively for early retirement (Kyyrä and Ollikainen, 2008).

³⁷ We discuss the robustness of our results with respect to this threshold value in section 3.5.3.

workforce by at least 50% between years t and $t + 1$ but were still in operation at the end of year $t + 1$ (*downsizing* subgroup).

These groups are followed over an 11-year period beginning three years before and ending seven years after the base year: from 1989 until 1999 and from 1994 until 2004, respectively. This results in two large unbalanced panel data sets. The 1992 sample has 2,471,751 observations (225,919 individuals) and the 1997 sample 2,872,552 observations (262,487 individuals).³⁸

We allow separations in the control group after the base year, implying that workers in the control group may be displaced later. In this respect, we follow Huttunen *et al.* (forthcoming), Hijzen *et al.* (2006), and Eliason and Storrie (2006). Some other studies require that individuals in the control group remain employed (possibly in the same firm) over the whole observation period. Of course, the members of the displacement and early-leaver groups can experience additional job losses in the later periods. Subsequent job losses can significantly increase the costs from the initial job loss (Stevens, 1997), whereas the likelihood of multiple job losses may be much higher during economic downturns (Eliason and Storrie, 2006).

3.4 Descriptive evidence

3.4.1 Macroeconomic environment

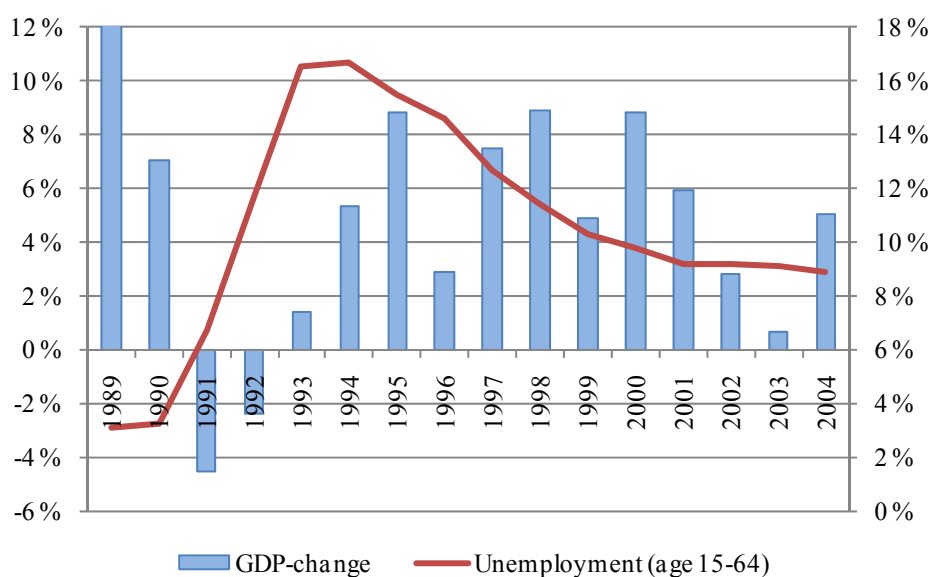
At the beginning of the 1990s, Finland suffered an exceptionally severe recession. GDP contracted three years in a row (1991–1993), and at the worst, in 1991 GDP decreased by over 6% (see Figure 3.1).³⁹ Overall, GDP declined by over 10% between 1991 and 1993, and the unemployment rate increased from 3.2% in 1990 to over 16% by 1993. According to one classification, episodes with peak-to-trough declines in output exceeding 10% are called depressions. Since the 1930s, there have been only a few such episodes in advanced economies before the ongoing global crisis triggered by the US subprime crisis in 2007. The latest depression occurred in Finland in the early 1990s. Hence, the experiences of Finnish workers who lost their jobs in 1992 represent an extreme case that highlights the consequences that may follow from a displacement during exceptionally difficult labour market conditions, providing a worst-case scenario for job losers.

³⁸ Persons disappear from our data only if they die or move abroad. In the 1992 sample the attrition is 3,608 persons from 1992 to 1999 and in the 1997 sample 3,706 persons from 1997 to 2004. There is no selection pattern according to the displacement status.

³⁹ For a discussion of the Finnish depression, see e.g. Honkapohja and Koskela (1999).

The depression years were followed by a period of strong and steady economic growth. At the end of the 1990s, the economy grew 3–6% a year and the unemployment rate declined about one percentage point each year. In 2001, economic growth slowed down and unemployment stabilized around the 9% level for the next few years. Compared to the depression sample, our second group displaced in 1997 encountered an entirely different situation. In many ways, this period is more typical and hence more in line with the research on displacement conducted elsewhere, although the unemployment rate remained at a high level by international standards.

Figure 3.1 GDP growth and unemployment rate in Finland, 1985–2004



Our depression period is related to the economic environment analyzed by Eliason and Storrie (2006). They followed Swedish workers displaced in 1987 until 1999. The displacements occurred during a boom period that was followed by a deep recession four years later, during which the unemployment rate rose from under 2% to over 8% and GDP fell by 6%. In other words, also Sweden experienced a severe recession in the early 1990s, though it was clearly less deep than the one experienced in Finland. Eliason and Storrie (2006) find that the negative impact of the displacement decreased over the first years but then started to increase at the beginning of the recession period. They concluded that displaced workers are more vulnerable to subsequent shocks, resulting in the long-lasting earnings losses for those whose post-displacement period is subject to adverse shocks. By implication, the prospects of earnings recovery are likely to be exceptionally poor for those people who lose their jobs in the middle of the recession period.

3.4.2 Background characteristics

We present summary statistics by displacement group status and sex in Table 3.1 and Table 3.2. Most of the variables are measured a year before the base period, but the earnings variables and plant and firm sizes are tracked for three pre-displacement periods. The outcome variable in our analysis is annual earnings, covering all salaries and wages received during a year. In section 5.3, we also discuss the results obtained using annual income as the outcome variable.

Table 3.1 Sample statistics for the 1992 sample

	Men					Women				
	Control	Displaced		Early-leavers		Control	Displaced		Early-leavers	
		d.s.	p.c.	d.s.	p.c.		d.s.	p.c.	d.s.	p.c.
Age	37.67	37.61	37.59	36.81	37.40	38.04	37.66	37.37	35.58	37.47
Years of education	11.36	11.08	11.03	11.71	10.88	11.06	10.87	10.88	11.13	10.76
Tenure (years)	11.30	9.06	8.72	7.27	7.97	11.25	10.79	9.78	8.30	8.74
Married (share)	0.66	0.63	0.65	0.60	0.61	0.61	0.58	0.61	0.59	0.58
Children under 7	0.40	0.37	0.42	0.42	0.42	0.23	0.27	0.29	0.48	0.30
Annual earnings:										
1989	22,814	21,106	21,551	22,213	21,025	15,811	14,869	15,062	14,423	15,197
1990	25,111	23,163	23,878	24,256	23,223	17,300	16,270	16,419	15,775	16,700
1991	26,169	23,727	24,244	24,701	23,191	18,424	16,707	16,478	15,251	16,743
Percentile:										
1989	64	62	64	59	62	45	41	45	35	43
1990	65	62	64	58	62	45	41	46	35	43
1991	65	61	63	56	60	45	38	42	30	40
Plant size:										
1989	467	145	65	294	73	364	188	49	286	64
1990	448	125	58	302	67	354	154	46	279	61
1991	436	111	41	266	50	339	140	39	271	45
Firm size:										
1989	2134	1259	645	973	751	2066	1588	381	1311	518
1990	1945	1178	523	954	527	1978	1650	324	1399	442
1991	1739	901	349	768	389	1757	1228	265	1012	334
Industry (share):										
Manufacturing	0.48	0.38	0.42	0.23	0.39	0.31	0.36	0.38	0.26	0.32
Construction	0.07	0.34	0.23	0.40	0.32	0.01	0.07	0.06	0.07	0.08
Trade	0.14	0.14	0.21	0.15	0.13	0.23	0.29	0.28	0.31	0.30
Transport	0.11	0.03	0.04	0.02	0.05	0.06	0.04	0.04	0.01	0.05
Business services.	0.10	0.08	0.08	0.14	0.07	0.25	0.21	0.22	0.23	0.18
Other	0.11	0.02	0.02	0.06	0.03	0.14	0.03	0.03	0.11	0.06
<i>N</i>	125,267	2175	1790	639	2728	89,961	1058	908	361	1032

Notes: Unless otherwise indicated, the numbers are for 1991. d.s. = downsizing subgroup. p.c. = plant-closure subgroup. Percentile = Earnings percentile within the plant. Trade also includes hotels and restaurants. Transport also includes telecommunications.

Table 3.2 Sample statistics for the 1997 sample

	Men					Women				
	Control	Displaced		Early-leavers		Control	Displaced		Early-leavers	
		d.s.	p.c.	d.s.	p.c.		d.s.	p.c.	d.s.	p.c.
Age	38.68	39.09	37.73	36.29	38.57	39.48	38.88	37.96	36.19	37.90
Years of education	11.40	11.30	11.63	11.71	11.12	11.29	11.45	11.25	11.70	11.07
Tenure (years)	12.52	12.63	8.52	8.54	12.20	12.06	11.21	9.01	9.07	10.27
Married (share)	0.61	0.60	0.57	0.54	0.57	0.57	0.56	0.52	0.55	0.55
Children under 7	0.42	0.38	0.50	0.46	0.37	0.24	0.28	0.23	0.46	0.33
Annual earnings:										
1989	25,881	26,817	25,289	24,576	24,377	19,091	18,655	17,602	18,481	17,779
1990	27,943	28,820	26,806	27,178	26,200	20,391	20,161	18,760	20,474	18,759
1991	29,240	30,295	27,497	28,931	28,041	21,403	20,774	19,714	19,738	19,770
Percentile:										
1989	61	58	67	56	54	42	40	45	33	44
1990	62	59	66	57	55	44	42	49	36	45
1991	63	60	66	60	56	45	42	49	32	45
Plant size:										
1989	370	371	43	261	134	304	346	44	435	99
1990	379	377	42	274	197	315	313	46	421	228
1991	399	533	36	276	124	329	363	40	442	134
Firm size:										
1989	1778	2370	898	797	894	1705	1541	743	1623	1041
1990	1813	2224	900	813	1194	1878	1507	761	1729	1697
1991	2160	1778	743	907	1232	1980	1351	749	1710	1704
Industry (share):										
Manufacturing	0.57	0.65	0.19	0.25	0.44	0.37	0.40	0.07	0.22	0.36
Construction	0.04	0.02	0.08	0.11	0.05	0.01	0.00	0.01	0.05	0.00
Trade	0.12	0.09	0.33	0.24	0.22	0.21	0.17	0.45	0.18	0.26
Transport	0.11	0.08	0.11	0.04	0.12	0.08	0.05	0.10	0.06	0.05
Business services.	0.09	0.09	0.19	0.29	0.17	0.19	0.16	0.30	0.45	0.30
Other	0.08	0.06	0.10	0.06	0.01	0.13	0.21	0.06	0.04	0.02
<i>N</i>	162,484	1075	336	112	327	96,724	820	248	108	253

Notes: Unless otherwise indicated, the numbers are for 1996. d.s. = downsizing subgroup. p.c. = plant-closure subgroup. Percentile = Earnings percentile within the plant. Trade also includes hotels and restaurants. Transport also includes telecommunications.

The displacement and control groups are similar in terms of age, education, and family background. With the exception of men displaced in 1997 from downsizing plants, the members of the displacement groups have somewhat shorter job tenures compared to the control group. The earnings percentile in the plant describes the worker's relative position in the earnings hierarchy within the employing plant. Since the average value of this measure is rather similar for displaced and control workers, there is no evidence of selective displacements towards the lower-paid employees within the plants.

As expected, the relative share of displaced workers is considerably higher in the 1992 sample compared to the 1997 sample: 2.6% vs. 0.9% (4.7% vs. 1.2% if the groups of early leavers are included).⁴⁰ The share of women seems to be slightly lower among workers displaced in 1992 and conversely higher in 1997 (33% vs. 43%), which is due to an exceptionally high layoff rate in the male-dominated construction sector in 1992. In the 1997 sample a disproportionate number of workers who lost their jobs in plant closures worked in trade (including also hotels and restaurants) and in business services, which are industries characterized by a high share of small business units. Not surprisingly, the average plant and firm sizes are smallest for the plant-closure subgroup.

There are only moderate differences between the early-leaver and displacement groups. With a few exceptions, the early leavers have more young children and have shorter job tenures on average. Differences in annual earnings are rather small, but the early leavers seem to be located at lower levels in the plant-specific earnings distributions, which may indicate that they did not quit but were dismissed. The average plant and firm sizes as well as the industry allocation are also quite different for the early leavers and displaced workers.

3.4.3 Empirical earnings distributions

Figure 3.2 shows the evolution of the 1st decile, median, and 9th decile of the annual earnings for the control and displacement groups of men displaced in 1992, Figure 3.3 describes the earnings for the men in 1997 group. The earnings development for women is described in Figure 3.4 and Figure 3.5. The earnings dispersion, as measured by the ratio of the 9th to 1st decile, increases quite strongly over time in the control group in all graphs. This trend is mainly driven by increasing variation in the annual working time within the control group. The declining pattern of the 1st decile, in particular, is due to an increasing fraction of control workers who leave full-time employment.

In the 1992 sample, the dip in earnings after displacement is very pronounced and clear at all parts of the distribution, which results from the substantial increase in non-employment. Median earnings declined from 1991 to 1993 by some 60% among displaced women and men. While the 9th decile declines less, the 1st decile drops all the way to zero and remains there until 1999 for the 1992 displacement groups, suggesting that a large share of displaced workers was out of work in each post-displacement year. Largely, the patterns of the earnings distributions of the two distinct displacement groups in the 1992 sample are very

⁴⁰ The total numbers of separations is, of course, much larger: Of those having been continuously employed in the same plant for the past three years, 18.5% and 12.3% separated during 1992 and 1997, respectively.

similar. The only notable difference is the higher 9th decile of male job losers in the plant-closure group over the pre-displacement years. Furthermore, the overall picture is very similar between women and men who were displaced in 1992.

Figure 3.2 The 1st decile, median, and 9th decile of (nominal) annual earnings for control and displacement groups for men displaced in 1992




Figure 3.3 The 1st decile, median, and 9th decile of (nominal) annual earnings for control and displacement groups for men displaced in 1997




Figure 3.4 The 1st decile, median, and 9th decile of (nominal) annual earnings for control and displacement groups for women displaced in 1992





Figure 3.5 The 1st decile, median, and 9th decile of (nominal) annual earnings for control and displacement groups for women displaced in 1997



The earnings quantiles of the displacement groups in the 1997 sample exhibit much smaller declines compared to the 1992 sample. There is no difference in the median earnings between those displaced in mass layoffs and the control group. Workers displaced from plants exiting in 1997 have lower median earnings in all years, and the difference increases somewhat after the displacement. Interestingly, the lower and upper tails of the distributions of the two displacement groups evolve somewhat differently over the post-displacement years. The 9th decile of the job losers of downsizing plants drops

only a little, while the 1st decile declines more clearly compared to the control group. In the plant-closure group, the 9th decile of the distribution follows very closely that of the control group, but the 1st decile of men declines sharply compared to the other displacement group. By contrast, there are no notable differences in the behaviour of the lower tail of the earnings distribution between women in the two displacement groups. As a result, displacement led to the largest increase in the earnings dispersion for male job losers in the plant-closure group.

Unlike in the 1992 sample, the earnings distribution of the plant-closure group in the 1997 sample differs from those of the other groups already in the pre-displacement years. This suggests the possibility that the plants exiting the market in the recovery period (and, hence, their employees) form a rather selective group. There may also be more early leavers in the 1997 sample because the chances of finding work in other plants were much better and because the economic environment was more predictable at the end of the 1990s. By conditioning on the control variables, we can eliminate these differences for men but not for women (see our quantile regression results below). Therefore, the female plant-closure group in the 1997 sample remains a problematic group due to a potential selection problem.

3.5 Quantile displacement effects

The descriptive analysis suggests that displacement had a notable and long-lasting negative effect on the earnings distribution. This effect seems clearly heterogeneous, as the ratio of the earnings quantiles between the displacement and control group was found to evolve differently over time at the lower and upper end of the distribution. We also recovered some differences in the background characteristics between the groups. Using the quantile regression method, we can model heterogeneity in the displacement effect while controlling for differences in the background characteristics.

Our specification for the conditional earnings quantiles differs slightly from the stylized example in section 3.2.2. First, we model relative effects by taking the log of the strictly positive values of annual earnings (but we do not drop zero earnings). Because of the equivariance of the quantiles to monotone transformations, this transformation of the dependent variable is completely transparent. Second, we have four groups of separating workers to be compared to the control group. These are indicated by the following dummy variables: d_{C_t} equals 1 for workers displaced in year t from plants that closed down in that year (plant-closure displacement group), d_{D_t} equals 1 for workers displaced in year t from downsizing plants (downsizing displacement group), e_{C_t} equals 1 for those who left their jobs in year t in plants that closed down in year $t + 1$ (plant-closure early-leaver group), and e_{D_t} equals 1 for those who left their jobs in year t in

plants that downsized in year $t + 1$ (downsizing early-leaver group). The control group – those who did not change the plant during year t – serves as the reference group in the analysis. Thus, our model for the conditional θ -th quantile of the earnings in year s is

$$F_{Y_s^*}^{-1}(\theta | \mathbf{z}_s) = \max \left\{ 0, \mathbf{x}_s' \boldsymbol{\beta}_{\theta s} + \alpha_{\theta s}^C d_{C_t} + \alpha_{\theta s}^D d_{D_t} + \eta_{\theta s}^C e_{C_t} + \eta_{\theta s}^D e_{D_t} \right\}, \quad (19)$$

where Y_s^* is log annual earnings for strictly positive earnings and zero otherwise, $t \in \{1992, 1997\}$ is the base period, $s \in \{t-3, t-2, \dots, t+7\}$, $\theta \in \{.1, .2, \dots, .9\}$, and $\mathbf{z}_s = (\mathbf{x}_s, d_{C_t}, d_{D_t}, e_{C_t}, e_{D_t})$. The vector of control variables \mathbf{x}_s includes age, age squared, pre-displacement tenure, education level (5 levels), place of residence (5 regions), marital status, indicator of children under the schooling age, the log annual earnings in year $t-4$ and the size category (4 classes) and industry (6 main industries) of the firm in year $t-1$. The past earnings are included to control for the effect of unobserved characteristics.

By taking the exponent of the right-hand side of (19), provided it is not zero, we obtain the conditional θ -th quantile of the annual earnings in year s . Coefficients of the group dummies capture proportional differences compared to the non-displacement case. For example, provided that $\mathbf{x}_s' \boldsymbol{\beta}_{\theta s} > 0$ and $\mathbf{x}_s' \boldsymbol{\beta}_{\theta s} + \alpha_{\theta s}^C > 0$, $\exp(\alpha_{\theta s}^C)$ gives the ratio of the θ -th quantile in year s if displaced in year t due to a plant closure to the θ -th quantile without displacement in year t . This proportional effect is independent of the values of control variables. If $\mathbf{x}_s' \boldsymbol{\beta}_{\theta s} > 0$ and $\mathbf{x}_s' \boldsymbol{\beta}_{\theta s} + \alpha_{\theta s}^C \leq 0$, the conditional θ -th quantile of annual earnings is zero with displacement but strictly positive without displacement, and thereby the ratio of the quantiles with and without displacement is zero, not $\exp(\alpha_{\theta s}^C)$. In other words, the proportional effect interpretation does not apply to arbitrarily values of control variables. This is a relevant concern when $\exp(\alpha_{\theta s}^C)$ is close to zero, which is the case with some lowest deciles in our application below. In those cases, the θ -th quantile of annual earnings of some people is predicted to drop to zero after displacement, which should be kept in mind when interpreting our results below.⁴¹

Using Powell's (1986) method,⁴² we have estimated the model (19) separately for women and men in the two samples. All regression parameters of each nine

⁴¹ If $\mathbf{x}_s' \boldsymbol{\beta}_{\theta s} \leq 0$ and $\mathbf{x}_s' \boldsymbol{\beta}_{\theta s} + \alpha_{\theta s}^C \leq 0$, the displacement has no effect at all because the conditional θ -th quantile is zero in any case. This is not a very relevant case in our application.

⁴² For details, see Jolliffe *et al.* (2000).

deciles were allowed to vary freely across the 11 cross sections. This amounted to 396 distinct quantile regressions. The point estimates of $\alpha_{\theta_s}^C$ and $\alpha_{\theta_s}^D$ are reported in Table 3.3 and Table 3.4, whereas the time patterns of $\exp(\alpha_{\theta_s}^C)$ and $\exp(\alpha_{\theta_s}^D)$ are shown in Figure 3.6–Figure 3.13. Each curve in the graphs shows how the proportional displacement effect at a particular decile evolves over time. In section 3.5.3 we also discuss briefly the results for the early leavers; in other words, the estimates of $\eta_{\theta_s}^C$ and $\eta_{\theta_s}^D$.

Table 3.3 Quantile displacement effects for 1992

θ	1989	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999
A. Men displaced from a downsizing plant, $\alpha_{\theta_s}^D$											
.1	-0.006	-0.014	-0.024	-0.851	-6.815	-6.806	-6.914	-6.100	-6.039	-5.851	-6.197
.2	-0.001	-0.010	-0.025	-0.500	-5.179	-3.404	-2.524	-1.896	-1.686	-1.192	-1.153
.3	-0.001	-0.008	-0.025	-0.330	-2.260	-1.723	-1.261	-0.964	-0.660	-0.511	-0.440
.4	0.000	-0.005	-0.023	-0.215	-1.410	-1.147	-0.719	-0.506	-0.359	-0.272	-0.229
.5	-0.002	-0.008	-0.020	-0.166	-0.834	-0.694	-0.402	-0.305	-0.227	-0.180	-0.165
.6	-0.004	-0.005	-0.022	-0.132	-0.479	-0.377	-0.272	-0.219	-0.173	-0.141	-0.132
.7	-0.003	-0.006	-0.020	-0.114	-0.283	-0.252	-0.203	-0.171	-0.147	-0.120	-0.113
.8	-0.003	-0.010	-0.020	-0.096	-0.225	-0.195	-0.161	-0.137	-0.126	-0.102	-0.099
.9	-0.008	-0.011	-0.027	-0.087	-0.171	-0.160	-0.124	-0.107	-0.097	-0.087	-0.091
B. Men displaced due to plant closure, $\alpha_{\theta_s}^C$											
.1	-0.007	-0.017	-0.028	-1.059	-7.622	-6.358	-5.773	-7.372	-8.671	-5.732	-5.778
.2	-0.004	-0.011	-0.015	-0.586	-7.252	-3.251	-3.093	-2.472	-1.468	-1.353	-1.154
.3	-0.005	-0.008	-0.015	-0.317	-2.437	-1.518	-1.278	-1.083	-0.690	-0.461	-0.460
.4	-0.005	-0.008	-0.014	-0.211	-1.528	-0.996	-0.660	-0.510	-0.349	-0.240	-0.227
.5	-0.001	-0.002	-0.009	-0.147	-0.916	-0.591	-0.349	-0.283	-0.214	-0.170	-0.173
.6	0.000	0.004	-0.014	-0.120	-0.422	-0.317	-0.228	-0.196	-0.155	-0.139	-0.141
.7	0.002	0.002	-0.014	-0.095	-0.250	-0.228	-0.177	-0.161	-0.138	-0.131	-0.133
.8	0.003	0.008	-0.016	-0.078	-0.210	-0.180	-0.150	-0.143	-0.131	-0.115	-0.126
.9	0.011	0.020	0.005	-0.077	-0.186	-0.148	-0.143	-0.124	-0.129	-0.096	-0.104
C. Women displaced from a downsizing plant, $\alpha_{\theta_s}^D$											
.1	-0.002	-0.006	-0.057	-1.056	-6.277	-5.641	-10.165	-5.582	-6.203	-5.253	-2.976
.2	-0.006	-0.005	-0.033	-0.651	-6.690	-4.613	-3.019	-1.989	-2.076	-2.064	-2.245
.3	-0.008	-0.004	-0.030	-0.428	-2.563	-1.825	-1.386	-0.969	-0.932	-0.915	-0.806
.4	-0.008	-0.009	-0.024	-0.245	-1.665	-1.264	-0.895	-0.647	-0.545	-0.432	-0.419
.5	-0.009	-0.009	-0.022	-0.153	-1.017	-0.882	-0.478	-0.353	-0.281	-0.248	-0.247
.6	-0.010	-0.012	-0.024	-0.106	-0.475	-0.459	-0.290	-0.228	-0.194	-0.169	-0.151
.7	-0.009	-0.012	-0.029	-0.090	-0.228	-0.227	-0.195	-0.167	-0.156	-0.128	-0.104
.8	-0.009	-0.012	-0.032	-0.068	-0.154	-0.173	-0.138	-0.133	-0.112	-0.107	-0.081
.9	0.008	0.003	-0.022	-0.046	-0.093	-0.110	-0.115	-0.085	-0.061	-0.036	-0.031
D. Women displaced due to plant closure, $\alpha_{\theta_s}^C$											
.1	-0.004	-0.011	-0.084	-1.282	-7.518	-5.599	-4.873	-4.820	-8.401	-3.239	-4.000
.2	0.002	-0.004	-0.059	-0.837	-7.544	-3.309	-3.151	-2.826	-2.174	-1.997	-1.396
.3	0.003	-0.003	-0.041	-0.450	-2.538	-1.764	-1.459	-1.237	-0.996	-0.923	-0.789
.4	0.002	-0.005	-0.035	-0.275	-1.587	-1.242	-0.919	-0.859	-0.655	-0.391	-0.337
.5	0.000	-0.007	-0.022	-0.188	-1.118	-0.865	-0.517	-0.448	-0.327	-0.253	-0.203
.6	-0.003	0.000	-0.025	-0.137	-0.568	-0.510	-0.281	-0.236	-0.219	-0.175	-0.161
.7	0.002	0.012	-0.024	-0.106	-0.258	-0.233	-0.183	-0.162	-0.158	-0.127	-0.125
.8	0.007	0.011	-0.025	-0.106	-0.191	-0.180	-0.146	-0.141	-0.135	-0.126	-0.121
.9	0.037	0.018	0.003	-0.043	-0.158	-0.139	-0.137	-0.129	-0.108	-0.112	-0.117

Notes: Significantly (95%-confidence level) non-zero coefficients in bold. Statistical inference based on the standard errors bootstrapped using 100 replications.

Table 3.4 Quantile displacement effects for 1997

θ	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004
A. Men displaced from a downsizing plant, $\alpha_{\theta_s}^D$											
.1	0.005	0.006	0.006	-0.088	-0.370	-0.465	-0.444	-0.278	-0.439	-0.768	-0.802
.2	0.006	0.006	0.006	-0.017	-0.075	-0.082	-0.072	-0.064	-0.063	-0.074	-0.076
.3	0.004	0.004	0.006	-0.005	-0.057	-0.045	-0.028	-0.023	-0.027	-0.021	-0.018
.4	-0.001	0.007	0.005	0.003	-0.042	-0.033	-0.020	-0.011	-0.020	-0.008	-0.004
.5	0.001	0.008	0.005	0.006	-0.036	-0.022	-0.018	-0.010	-0.019	-0.008	-0.012
.6	0.004	0.010	0.006	0.006	-0.024	-0.014	-0.020	-0.004	-0.004	-0.010	-0.016
.7	0.006	0.020	0.006	0.019	0.001	-0.006	-0.015	0.012	0.001	-0.006	-0.002
.8	0.006	0.024	0.013	0.022	0.020	0.006	-0.010	0.010	-0.004	-0.007	-0.012
.9	0.004	0.024	0.018	0.031	0.023	0.004	-0.016	0.004	-0.003	-0.003	0.003
B. Men displaced due to plant closure, $\alpha_{\theta_s}^C$											
.1	-0.031	-0.023	-0.057	-0.412	-1.042	-1.138	-1.164	-2.005	-1.558	-1.793	-1.769
.2	-0.012	-0.010	-0.038	-0.125	-0.285	-0.205	-0.291	-0.215	-0.280	-0.252	-0.288
.3	-0.007	-0.003	-0.026	-0.087	-0.149	-0.100	-0.137	-0.091	-0.136	-0.132	-0.137
.4	-0.006	-0.007	-0.014	-0.057	-0.095	-0.051	-0.091	-0.063	-0.078	-0.086	-0.100
.5	-0.005	-0.002	-0.013	-0.026	-0.052	-0.034	-0.049	-0.013	-0.056	-0.039	-0.055
.6	0.003	0.003	-0.011	-0.028	-0.024	-0.016	-0.029	-0.009	-0.025	-0.032	-0.031
.7	0.004	-0.003	-0.020	-0.014	-0.020	-0.008	-0.030	-0.024	-0.019	-0.034	-0.048
.8	-0.010	-0.024	-0.027	-0.027	-0.030	-0.008	-0.013	-0.027	-0.011	-0.041	-0.038
.9	0.011	-0.005	-0.045	-0.063	-0.027	0.003	-0.013	-0.033	-0.040	-0.050	-0.058
C. Women displaced from a downsizing plant, $\alpha_{\theta_s}^D$											
.1	-0.010	-0.005	-0.019	-0.524	-1.446	-1.815	-1.554	-0.901	-0.868	-2.247	-2.904
.2	-0.005	-0.003	-0.003	-0.187	-0.487	-0.318	-0.256	-0.213	-0.248	-0.365	-0.421
.3	-0.009	-0.002	-0.005	-0.056	-0.087	-0.075	-0.064	-0.089	-0.092	-0.138	-0.134
.4	-0.009	-0.008	-0.009	-0.028	-0.062	-0.057	-0.052	-0.066	-0.061	-0.078	-0.075
.5	-0.011	-0.014	-0.016	-0.027	-0.057	-0.052	-0.048	-0.056	-0.053	-0.066	-0.080
.6	-0.011	-0.020	-0.019	-0.032	-0.064	-0.061	-0.050	-0.053	-0.053	-0.057	-0.083
.7	-0.026	-0.034	-0.036	-0.047	-0.075	-0.078	-0.063	-0.066	-0.063	-0.061	-0.073
.8	-0.043	-0.047	-0.050	-0.057	-0.086	-0.091	-0.082	-0.086	-0.083	-0.085	-0.083
.9	-0.061	-0.062	-0.061	-0.062	-0.101	-0.104	-0.101	-0.090	-0.102	-0.084	-0.099
D. Women displaced due to plant closure, $\alpha_{\theta_s}^C$											
.1	-0.055	-0.104	-0.080	-0.697	-2.141	-3.139	-2.687	-1.142	-1.320	-1.531	-4.320
.2	-0.031	-0.023	-0.022	-0.239	-0.640	-0.868	-0.626	-0.412	-0.303	-0.259	-0.312
.3	-0.011	-0.007	-0.010	-0.148	-0.353	-0.292	-0.329	-0.250	-0.226	-0.205	-0.171
.4	0.002	-0.003	-0.015	-0.074	-0.194	-0.189	-0.190	-0.199	-0.138	-0.123	-0.120
.5	0.003	0.006	-0.002	-0.025	-0.132	-0.131	-0.100	-0.137	-0.091	-0.089	-0.087
.6	0.008	0.007	-0.005	-0.021	-0.076	-0.101	-0.069	-0.103	-0.048	-0.051	-0.042
.7	0.014	-0.009	-0.003	-0.007	-0.078	-0.067	-0.068	-0.076	-0.017	-0.004	-0.040
.8	0.017	-0.020	0.003	0.015	-0.048	-0.044	0.007	-0.025	-0.021	-0.005	0.014
.9	0.046	0.017	0.038	0.042	-0.015	0.026	0.058	0.077	0.050	0.059	-0.008

Notes: Significantly (95%-confidence level) non-zero coefficients in **bold**. Statistical inference based on the standard errors bootstrapped using 100 replications.

Figure 3.6 Men displaced from downsizing plants in 1992, proportional quantile displacement effects $\exp(\alpha_{\theta_s}^D)$




Figure 3.7 Men displaced from downsizing plants in 1997, proportional quantile displacement effects $\exp(\alpha_{\theta_s}^D)$




Figure 3.8 Men displaced from closing plants in 1992, proportional quantile displacement effects $\exp(\alpha_{\theta_s}^C)$




Figure 3.9 Men displaced from closing plants in 1997, proportional quantile displacement effects $\exp(\alpha_{\theta_s}^C)$




Figure 3.10 Women displaced from downsizing plants in 1992, proportional quantile displacement effects $\exp(\alpha_{\theta_s}^D)$




Figure 3.11 Women displaced from downsizing plants in 1997, proportional quantile displacement effects $\exp(\alpha_{\theta_s}^D)$




Figure 3.12 Women displaced from closing plants in 1992, proportional quantile displacement effects $\exp(\alpha_{\theta_s}^C)$





Figure 3.13 Women displaced from closing plants in 1997, proportional quantile displacement effects $\exp(\alpha_{\theta_s}^C)$



3.5.1 Pre-displacement effects

Under the assumption that the displacement status is exogenous, we can interpret differences in the conditional earnings distributions between the control and displacement groups as the causal effect of displacement. This conditional independence assumption also implies that, given the control variables, there should be no notable earnings differences between the groups in the periods when the displacement group was not yet affected. Although the earnings of the displacement group may have been affected some time before the actual displacement took place, the earnings differences between the groups should disappear at some point when we go further back in time. If that does not happen, we take it as evidence against the validity of the conditional independence assumption.

As seen in Table 3.3, the earnings distributions of workers who were displaced in 1992 due to plant downsizing or closure are very similar to that of the control group three years before the displacement year. Only one displacement dummy out of 18 for men in Table 3.3 gets a significantly (at the 5% risk level) non-zero coefficient in 1989. Namely, the 1st decile of the earnings distribution of men losing their jobs due to plant closure in 1992 is estimated to be 0.7% lower than that of the control group. In Table 3.3, there are five statistically significant coefficients for women in 1989. Except for the highest decile for women in the plant-closure group, all significant effects imply less than 1% difference in the deciles compared to the control group. These very small discrepancies between the female groups may result from sample noise, because only one coefficient differs significantly from zero a year later in 1990 (the 4th decile in panel C in Table 3.3). It should be stressed that at the 1% risk level, which might be a more reasonable choice given our sample sizes, only one coefficient out of 36 differs significantly from zero in 1989 in Table 3.3.⁴³ In other words, the earnings distributions of the displacement and control groups for both sexes are virtually identical four years before the base period, suggesting that the conditional independence assumption holds in the 1992 sample.

There are some statistically significant differences for men in 1990, but these are very modest in absolute value. For both women and men we find a decline in the earnings distributions of the displacement groups in 1991. At the left tail of women's distribution, the difference compared to the control group exceeds 5%, but elsewhere the differences are still only around 2%. This implies a possibility that the annual working hours of female employees with relatively low earnings have a tendency to decline prior to job loss. Many of these workers are probably

⁴³ Because of the large number of point estimates, we are expected to recover some spuriously significant effects even if there were no true effects.

part-timers whose working time and, hence, earnings vary with firm-specific business conditions. This could explain a stronger effect for women who are more likely to be employed on a part-time basis.

Overall, the estimates for the pre-displacement periods show very modest earnings differences between the displacement and control groups in the 1992 sample. One might expect these small pre-displacement effects to be attributable to an unexpected nature of the 1990s depression. This explanation does not sound very convincing because the pre-displacement effects appear to be rather small, or even smaller, in the 1997 sample (see Figure 3.7, Figure 3.9 and Table 3.4). In the 1997 sample, the earnings distributions of the displacement and control groups in 1994 are equally similar for men as they were in 1989 in the 1992 sample. In particular, none of the effects for men in 1994 differs statistically significantly from zero in Table 3.4, which is consistent with the conditional independence assumption. Surprisingly, one year later the upper half of the earnings distribution of men to be displaced from downsizing plants lies slightly *above* that of the control group. In addition to being very small in absolute value, these differences also vanish in the next year (with the exception of the effect at the 9th decile). The displacement effect for men in the plant-closure group becomes statistically significant for the first time in 1996 at the lower end of the distribution.

For women our setting does not work quite as well. The group of women to be displaced from downsizing plants seem to earn less than the control group already in 1994 (panel C in Table 3.4). The differences in the deciles are relatively small, being 0.9–5.9%, but statistically significant from the 3rd decile upwards and around 5% at the right tail of the distribution. This raises a doubt that our sample design fails and a selection problem remains for this particular group of women. For women in the plant-closure group we find only one statistically significant coefficient in 1994. The absolute value of this effect is relatively high, implying a difference of 5.4% in the 1st decile compared to the control group. On the other hand, the effect is associated with a rather high standard error (0.020) and it disappears by 1996.

To summarize, with the exception of women displaced in 1997 due to plant downsizing, the earnings distributions of all other displacement groups are almost identical to that of the control group three or four years prior to the base period. Hence, our approach to detect exogenous displacements using mass layoffs and plant closures seems successful, though one female group might pose some problems. We also found relatively small pre-displacement effects one to two years before the displacement year.

3.5.2 Post-displacement effects

Losing a job in the middle of the depression in 1992 has a huge effect on the earnings distribution. In the year following the displacement, the median earnings of displaced men are less than half of the median of the control group (see Figure 3.6 and Figure 3.8). The decline in the median earnings is even slightly more pronounced for displaced women. Not only does the entire distribution shift down, but also its shape changes drastically. The deciles above the median drop less, whereas the lower deciles decline much more than the median. The huge proportional effect in the left tail of the distribution implies that a large fraction of displaced workers were unable to return to paid employment

After the sharp initial drop in the first two years following the job loss, the earnings distributions of the displacement groups start to converge toward that of the control group. This recovery is rather strong between the 2nd and 6th deciles, though it slows down after a few years for women who lost their jobs in mass layoffs. This pattern can be attributed to an increasing probability of having found a suitable job after displacement.⁴⁴ It is striking, however, that the 1st decile does not show any sign of recovery, suggesting that the displaced worker has a notable risk of remaining outside paid work at the end of the follow-up period.

Displacement has a much weaker effect in the upper part of the distribution. At the 9th decile, for example, the displacement effect never exceeds 20%. Hence, a job loser can sometimes perform relatively well in the labour market compared to what would have happened without the displacement. On the other hand, the recovery of the upper deciles is rather slow and, consequently, the upper tail does not catch up with the counterfactual level by the end of the observation period. Except for the effect at the 1st and 9th deciles for women displaced from downsizing plants, the displacement effect is statistically significant at all deciles in 1999 (see Table 3.3). The two lowest deciles for the displaced worker are less than half of the counterfactual values without displacement, whereas the medians and upper deciles are 3–25% below the counterfactual values. It is remarkable that the displacement in the depression still has an effect on the entire distribution seven years after the job loss. Note also that the consequences of job loss in 1992 are rather similar for workers displaced from downsizing plants and those displaced from exiting plants.

⁴⁴ The average number of employment months in the displacement group increased quite rapidly from 1993 to 1998.

Women are subject to larger long-term earnings losses than men. Among those displaced from downsizing plants, the displacement effect at the median is 15% and 22% in 1999 for men and women, respectively. Except for the two lowest deciles, the sex difference in the displacement effects is even larger below the median, whereas it is rather small in the upper part of the distribution. A similar pattern exists for the plant-closure group, though the sex difference in the displacement effects at the median is less pronounced for that group.

When we turn our attention from the severe depression to the ensuing recovery period, the picture of earnings losses changes dramatically (see Figure 3.7, Figure 3.9, Figure 3.11 and Figure 3.13). Compared to the 1992 effect, the effect of displacement in 1997 is much smaller on average but even more concentrated in the lower end of the distribution. The displacement effect exhibits relatively little variation between the 3rd and 9th deciles (compared to the 1992 sample), but it is of a different order of magnitude at the two lowest deciles. In all displacement groups, the lowest decile drops more than 50% compared to the control group. Despite very large absolute values, the displacement effect at the 1st decile is occasionally statistically insignificant in the last years of the observation period (see Table 3.4), which may be due to the relatively small sample sizes of the 1997 displacement groups.

In 1997, the displacement has a statistically significant negative effect for men only at the lower end of the earnings distribution. This effect is more pronounced for those who were displaced from exiting plants. The displacement effect in the lower part of the earnings distribution gets stronger in 1998, the median being around 5% below the median of the control group. The displacement effect for men is not statistically significant above the 6th decile after 1997 except at one decile in 2004 for the plant-closure displacement group. By the end of the observation period, the displacement effect for men at the median disappears but the negative effect at the lower deciles remains statistically significant. This heterogeneity implies that a large fraction of workers experienced rather moderate earnings losses, if any, after a job displacement that took place during the period of economic growth. In other words, much of the decline in mean earnings can be attributed to a relatively small group of displaced workers who suffered from notable earnings reductions due to difficulties to return to work.

Women displaced in 1997 due to plant closures suffered from larger earnings losses than male job losers. One year after the displacement, the effect is statistically significant between the 1st and 8th deciles for the plant-closure group (see panel D in Table 3.4). This displacement effect gets weaker over time, but remains statistically significant in the lower half of the distribution until the end of the observation period. In 2004, the effect is still some 8% at the median and even larger at the lower deciles. The effect of being displaced from an exiting plant is slightly larger in 2004 for women than for men. In the earlier years, the displacement effect is clearly stronger for women. These findings are in line with

the larger losses for women in the 1992 sample. There seems to be a positive displacement effect for the female plant-closure group at the highest decile over the years 1999–2003, but this effect is not significantly different from zero in any year.

The displacement effect for women who lost their jobs in mass layoffs is significantly negative at each decile between 1997 and 2004 (panel C in Table 3.4). In other words, the entire earnings distribution of this group is estimated to lie below the distribution of the control groups for the complete post-displacement period. The effect of being displaced from a downsizing plant for women is relatively weak but very persistent, as there is little or no recovery at all over time. These estimates should be treated with caution, however, as the pre-displacement effects raised some doubts about the validity of the conditional independence assumption for this particular group.

In the 1992 sample, there are no systematic differences in the displacement effects between workers who lost their jobs in mass layoffs and those who became displaced from exiting plants. By contrast, men in plants that closed down in 1997 were subject to much larger earnings reduction than those displaced from downsizing plants in that year. Our findings for women who were displaced in 1997 are inconclusive due to a potential selection problem in the downsizing subgroup. A potential explanation is that, during the exceptionally deep (and to some extent unexpected) depression, mass layoffs and plant closures were rather common, and hence perhaps less selective, events.

One of our key findings is that the effect of displacement is very heterogeneous, being much larger in the lower quantiles, and this holds for women and men in both periods. This also implies a higher degree of earnings dispersion, as measured by the ratio of the upper deciles to the lower ones, for the displaced workers. In other words, job loss increases uncertainty about the future earnings level, suggesting an additional welfare loss for the risk-averse workers.

3.5.3 Robustness of the results

We have checked the robustness of the main results with respect to various departures from our benchmark setting. Here we describe the main findings briefly, but do not report any parameter estimates due to a huge number of them.⁴⁵ First, our analysis involves an implicit assumption that better workers did not quit and less able workers were not laid off in any significant scale in the periods preceding a mass layoff or plant closure. In section 3.5.1 we did not find evidence of notable differences in the earnings distributions between the

⁴⁵ Of course, the detailed results are available from the authors on request.

treatment and control groups 3–4 years before the base periods (except for women in the 1997 downsizing group). While these findings support the validity of our sample design, it is of interest to compare the earnings of the early leavers to that of the control and displacement groups. As discussed earlier, two subgroups of early leavers are included in the data: workers who separated in the base year (1992 or 1997) from plants that downsized or closed down during the *next* year. Differences in their earnings distributions compared to that of the control group are captured by the coefficients $\eta_{\theta_s}^C$ and $\eta_{\theta_s}^D$ in (19).

In the 1992 sample, these effects exhibit very similar patterns over time but are (almost) uniformly larger (in absolute value) than the associated displacement effects, $\alpha_{\theta_s}^C$ and $\alpha_{\theta_s}^D$, in the post-displacement period. In other words, between 1992 and 1999 each decile of the earnings distribution of the early-leaver group typically dropped more than the corresponding decile of the associated displacement group. This result is consistent with the hypothesis that the employers laid off their least productive employees prior to the period of mass layoff or plant closure. On the other hand, we do not detect notable differences in the earnings distribution between the early leavers and the control groups 3–4 years before 1992,⁴⁶ implying that the early leavers are *not* a selective group compared to the control and displacement groups. Following Gibbons and Katz (1991), a potential explanation is that the early leavers are otherwise similar but are affected by a stigma effect compared to those who separated during the period of a mass layoff or plant closure.

In the 1997 sample, we find much larger earnings losses for the early-leaver group than for the displacement group. In this case, there seems to be a selection problem. With an exception of men in downsizing plants, the earnings distribution of the early leavers was below that of the control group already in 1994. Therefore, the early leavers in 1997 seem to be a selective group of workers in terms of unobserved characteristics, and thereby their exclusion from the control and displacement groups is important for appropriate statistical inference.

Second, our threshold value for mass layoffs – a 50% reduction in employment – is essentially arbitrary. We have checked the robustness of our results with respect to this choice by lowering the threshold value to 30%. Our results remain qualitatively unchanged, but the displacement effect right after the period of job loss becomes a few percentage points stronger at the 6th and lower deciles in the 1992 group and at the 1st and 2nd decile in the 1997 group. We interpret this as an effect of including more selective dismissals in our displacement group.

⁴⁶ At the 5% risk level, 3 out of 36 coefficients differ significantly from zero in 1989.

Third, so far we have discussed the consequences of job losses that occurred at two specific points in time. Those years were not randomly chosen but we have conducted similar analyses for all displacements taking place in the period 1992–2001. The results from this exercise show that workers displaced in 1992 were subject to the largest earnings losses. The earnings losses exhibit a gradually decreasing trend as a function of the displacement year from 1992 to 1997. There is no notable variation in the displacement costs with respect to the timing of job loss after 1997. Thus, our results for the 1997 sample describe the costs of job loss under “normal” economic conditions, whereas the 1992 results provide an upper bound for the displacement costs that the worker can face in an exceptionally difficult economic environment.

Finally, as the dependent variable in our analysis is earnings, our estimates describe a reduction in labour income that results from shorter working time and/or lower wage rates following the displacement. When out of work, individuals are typically entitled to income transfers, like unemployment benefits, disability benefits and/or housing allowance, which can compensate for a large part of the earning losses in a welfare state like Finland. Therefore, we should expect displaced workers to experience smaller income losses than earnings losses. To address this issue we have replicated our analysis by using taxable annual income (excluding capital income) as the dependent variable in place of labour income. As expected, the displacement effect on annual income is much smaller than on annual earnings at the lower end of the distribution. In the 1992 sample, the displacement effect at the two lowest deciles of annual income is above 30% in 1999. The corresponding effect at the lower end of the distribution for the 1997 sample in 2004 is smaller but statistically significant, being around 10%. At the upper end of the distributions, the effects of job loss on annual income and earnings are of the same magnitude. These findings are not very surprising given that eligibility for income transfers depends on the level of labour income. Nevertheless, it may come as a surprise that the displacement has an equally long-lasting effect on the distribution of annual income as it has on the earnings distribution. Namely, seven years later, a job loss still has a statistically significant effect on the entire distribution of annual income for those displaced in 1992 and at the lower end of the distribution for those displaced in 1997.

3.5.4 Comparisons to results from mean regressions

In Table 3.5 we report the “ordinary” fixed effects estimates of the expected earnings losses. We calculated these for illustrating the conventional displacement effect results and to make it easier to compare the results with evidence from other countries. In these regressions the annual earnings of worker i in year s , Y_{is} , is regressed on an individual fixed effect and the same set of displacement dummies and control variables as in equation (19), excluding the time invariant variables. We ran the regressions for the complete samples (i.e. for the same samples used in the quantile regressions) using Y_{is} as the dependent

variable, as well as for subsamples of observations with strictly positive earnings using both Y_{is} and $\log(Y_{is})$ as the dependent variable.

Not surprisingly, the mean earnings losses are much larger for those displaced in 1992 than those losing their jobs in 1997. When all observations are included in the analysis, the displacement effect estimates are uniformly higher. Removing observations with zero earnings from the analysis reduces the displacement effects by some 10–20% in the post-displacement years for men and somewhat more for women. Thus the use of the selective sample of observations with strictly positive earnings leads to underestimation of the displacements costs and to potentially misleading comparisons between women and men. There is no evidence of statistically significant mean displacement effects for workers displaced in 1997 due to plant closure (except for men with strictly positive earnings), even though we found significant displacement effects at the lower end of the earnings distribution in the quantile regression analysis. In other words, one may miss a relatively small and heterogeneous effect by looking only at the mean effects.

Table 3.5 Fixed effects estimates of mean earnings losses

Displacement from a downsizing plant													
	Displacement in 1992						Displacement in 1997						
	Men			Women			Men			Women			
	All	$Y_{is} > 0$	$\log(Y_{is})$	All	$Y_{is} > 0$	$\log(Y_{is})$	All	$Y_{is} > 0$	$\log(Y_{is})$	All	$Y_{is} > 0$	$\log(Y_{is})$	
1990	-17	-28	0.0079	0	2	0.0037	1995	198	212	-0.0037	275	261	0.0155
	(504)	(502)	(0.0108)	(255)	(240)	(0.0177)		(2,059)	(2,069)	(0.0121)	(536)	(529)	(0.0175)
1991	-505	-514	-0.0023	-661	-667	-0.0351	1996	398	410	-0.0118	-110	-129	-0.0095
	(504)	(502)	(0.0108)	(255)	(240)	(0.0177)		(2,059)	(2,069)	(0.0121)	(536)	(529)	(0.0175)
1992	-5,628	-5,520	-0.3133	-3,951	-3,750	-0.3489	1997	-635	-605	-0.0548	-1,046	-1,032	-0.1336
	(506)	(505)	(0.0109)	(255)	(242)	(0.0179)		(2,070)	(2,081)	(0.0121)	(536)	(530)	(0.0175)
1993	-12,800	-10,983	-1.0286	-8,783	-7,188	-0.9337	1998	-2,349	-1,723	-0.1369	-2,718	-2,002	-0.2210
	(507)	(546)	(0.0117)	(255)	(265)	(0.0196)		(2,077)	(2,113)	(0.0123)	(536)	(540)	(0.0178)
1994	-11,874	-10,587	-0.8749	-8,014	-6,781	-0.7683	1999	-3,930	-3,326	-0.1225	-2,315	-1,451	-0.1541
	(508)	(541)	(0.0116)	(255)	(260)	(0.0193)		(2,077)	(2,121)	(0.0124)	(536)	(545)	(0.0180)
1995	-10,793	-9,459	-0.6978	-7,008	-5,701	-0.5693	2000	-4,350	-3,790	-0.1078	-1,487	-861	-0.1105
	(508)	(539)	(0.0116)	(256)	(260)	(0.0192)		(2,076)	(2,120)	(0.0124)	(536)	(544)	(0.0180)
1996	-9,701	-8,340	-0.5366	-6,142	-5,217	-0.4927	2001	-2,897	-2,374	-0.0945	-1,265	-663	-0.0885
	(508)	(539)	(0.0116)	(256)	(256)	(0.0189)		(2,078)	(2,121)	(0.0124)	(536)	(544)	(0.0180)
1997	-8,808	-7,705	-0.4639	-5,317	-4,225	-0.4038	2002	-3,283	-2,792	-0.0869	-1,230	-813	-0.0931
	(509)	(537)	(0.0115)	(256)	(258)	(0.0191)		(2,081)	(2,133)	(0.0124)	(537)	(544)	(0.0179)
1998	-8,252	-7,096	-0.3661	-4,969	-3,770	-0.3253	2003	-3,100	-2,392	-0.0894	-1,574	-1,085	-0.1115
	(511)	(537)	(0.0115)	(256)	(258)	(0.0191)		(2,088)	(2,155)	(0.0126)	(537)	(547)	(0.0181)
1999	-7,600	-6,430	-0.3554	-2,291	-435	-0.2634	2004	-1,972	-1,154	-0.0580	-1,431	-992	-0.0816
	(511)	(537)	(0.0116)	(256)	(260)	(0.0192)		(2,089)	(2,166)	(0.0126)	(537)	(552)	(0.0182)
Displacement from a closing plant													
1990	317	314	0.0172	-77	-65	0.0054	1995	-655	-614	0.0006	-318	-331	0.0302
	(555)	(553)	(0.0119)	(275)	(259)	(0.0191)		(3,675)	(3,693)	(0.0215)	(971)	(960)	(0.0317)
1991	-420	-422	0.0001	-1,211	-1,201	-0.0698	1996	-1,354	-1,308	-0.0215	-345	-358	0.0319
	(555)	(553)	(0.0119)	(275)	(259)	(0.0191)		(3,675)	(3,693)	(0.0215)	(971)	(960)	(0.0317)
1992	-5,922	-5,620	-0.3400	-4,834	-4,569	-0.4263	1997	-3,562	-3,420	-0.1497	-2,287	-2,305	-0.1596
	(555)	(556)	(0.0119)	(275)	(261)	(0.0193)		(3,692)	(3,715)	(0.0217)	(972)	(961)	(0.0317)
1993	-13,382	-11,374	-1.0361	-9,589	-7,819	-0.9522	1998	-4,567	-3,608	-0.2127	-3,797	-3,061	-0.3165
	(556)	(603)	(0.0130)	(275)	(286)	(0.0211)		(3,706)	(3,795)	(0.0221)	(977)	(989)	(0.0326)
1994	-11,693	-10,118	-0.7985	-8,634	-7,552	-0.8337	1999	-5,079	-4,038	-0.1415	-3,719	-2,513	-0.1937
	(557)	(593)	(0.0128)	(275)	(278)	(0.0206)		(3,706)	(3,815)	(0.0222)	(977)	(1,003)	(0.0331)
1995	-10,909	-9,300	-0.6573	-7,609	-6,220	-0.5881	2000	-5,384	-4,538	-0.1435	-3,011	-2,087	-0.1382
	(557)	(594)	(0.0128)	(275)	(280)	(0.0207)		(3,710)	(3,815)	(0.0222)	(976)	(998)	(0.0329)
1996	-10,263	-8,817	-0.5752	-7,189	-5,891	-0.5285	2001	-3,444	-2,319	-0.1116	-2,119	-1,341	-0.0453
	(558)	(592)	(0.0127)	(276)	(280)	(0.0207)		(3,714)	(3,849)	(0.0224)	(976)	(996)	(0.0329)
1997	-9,202	-8,045	-0.4417	-6,436	-5,374	-0.4653	2002	-3,933	-2,795	-0.0874	-1,413	-644	0.0313
	(559)	(588)	(0.0126)	(276)	(278)	(0.0206)		(3,706)	(3,849)	(0.0224)	(979)	(1,000)	(0.0330)
1998	-8,829	-7,460	-0.3597	-6,043	-5,026	-0.4138	2003	-3,151	-2,218	-0.0721	-1,041	-548	0.0149
	(560)	(590)	(0.0127)	(276)	(277)	(0.0205)		(3,706)	(3,853)	(0.0225)	(979)	(1,001)	(0.0331)
1999	-9,702	-8,690	-0.3611	-5,581	-4,949	-0.3755	2004	-3,345	-2,540	-0.0876	-1,571	-1,071	0.0155
	(560)	(589)	(0.0127)	(276)	(275)	(0.0203)		(3,714)	(3,872)	(0.0226)	(980)	(1,012)	(0.0334)

Notes: Displacement effects from model $Y_{is} = a_i + \gamma_s + \mathbf{x}_{is}^v \beta + \alpha_s^C d_{isC} + \alpha_s^D d_{isD} + \eta_s^C e_{isC} + \eta_s^D e_{isD} + \varepsilon_{is}$ where a_i are the person fixed effects, γ_s are the year effects, $s \in \{t-3, t-2, \dots, t+7\}$ and $t \in \{1992, 1997\}$. The dummies for the displacement groups are the same as in equation (6). $Y_{is} > 0$ marks estimates where observations with no earnings are dropped. The displacement effects are set to 0 for two pre-displacement years, 1988 and 1989 for workers displaced in 1992, and 1993 and 1994 for workers displaced in 1997. Significantly (95%-confidence level) non-zero coefficients in **bold**. Standard errors in (parentheses).

3.6 Concluding remarks

We analysed the earnings losses owing to involuntary job loss among Finnish workers who became displaced during a period of depression or recovery. Using the quantile regression method, we estimated the effect of displacement at each decile of the earnings distribution. Our findings from both time periods suggest that 1) displaced workers suffer from substantial and persistent earnings losses, 2) women are subject to larger earnings losses than men, and 3) the effect of displacement is very heterogeneous, being much larger in the lower quantiles and implying an increase in earnings dispersion following displacement.

The first finding is in accordance with the results from the US labour market. The results from other European labour markets are mixed in this respect. The second result is interesting given that most US studies have not found notable differences between women and men (see e.g. table 1 in Couch and Placzek, 2010), whereas the European studies have not paid much attention to the gender aspect of displacement costs. In Finland, as in other countries, women are more frequently out of work for family reasons. That may induce employers to favour male employees when investing in managerial and professional skills. If so, and if such skills are generally transferable, i.e. not lost in job displacement, one can expect to find smaller earnings losses for displaced men. Kletzer (1989) finds that managerial, professional and technical workers retain a larger share of their returns to seniority after job loss than blue collar workers do. Taken together with the fact that managerial, professional and technical jobs are disproportionately held by men, employers' investment behaviour may lead to larger earnings losses for women. Still, it remains unclear why earnings losses differ between men and women in Finland, but not in the US. It should be stressed that women's labour market position is quite different in Finland. On the one hand, the relatively generous maternity and parental leave schemes encourage career breaks, but on the other hand, public day care and school meals help the mothers of young children to work full-time if they want to. Moreover, compared to most other countries, part-time work among women is not very common, whereas the labour force participation rate of women is rather high in Finland.

The third finding, the heterogeneity in the displacement effect, has important implications. First, the positive effect on earnings dispersion means that job loss does not only cause a significant decline in expected earnings, but also raises uncertainty about the level of future earnings. This suggests an additional welfare loss for risk-averse workers, the effect of which has been typically ignored in the discussion of displacement costs. Secondly, the dominant effect at the lower end of the distribution is consistent with the hypothesis that the relative importance of transferable individual-specific skills, which are not lost in job displacement, is larger for high-ability workers, who tend to populate the upper part of the conditional earnings distribution. Finally, the disproportionately large effect on

the first two deciles implies that the effect on the expected earnings loss is in large part driven by an increased risk of joblessness and low-paid employment following job displacement. This implies that job training and job replacement programmes targeted at unemployed job seekers, if effective in enhancing re-employment, can provide a means to reduce the average displacement cost.

By contrasting the results of the two periods, we found much larger earnings losses for those who lost their jobs during the depression period. Men (women) who were displaced in the middle of the depression had approximately 15% (20%) lower median earnings seven years after the job loss. Because of the exceptionally difficult labour market conditions, their earnings distribution as a whole remained below the counterfactual level until the end of the follow-up period. By contrast, job loss in the recovery period had a long-lasting effect only in the lower half of the distribution. For women displaced in 1997, the median effect seven years later was about 8%, whereas it did not differ from zero for men. These long-term losses do not vanish even when income transfers are accounted for. These results complement the findings of Eliason and Storrie (2006) and Couch and Placzek (2010) about the role of the business cycle in determining the size of earnings losses. Eliason and Storrie (2006) find that the displacement effect, which was first decreasing, started to increase when the Swedish economy was hit by a recession. By comparing their results from a different period in more favourable economic times to the results of Jacobson *et al.* (1993), Couch and Placzek (2010) conclude that the earnings losses in the US labour market are smaller after displacement in “ordinary” economic times than in an economic downturn.

Given that the world economy is currently experiencing its deepest downturn in the post-World War II period, our results obtained from the depression period are of particular interest. The extent of output and employment losses for many countries is projected to be of a similar magnitude that Finland experienced in the early 1990s, i.e. much bigger than those in the periods studied by Eliason and Storrie (2006) and Couch and Placzek (2010). For example, GDP declined in 2009 by 4.9% in the United Kingdom, 5.0% in Germany, 7.1% in Ireland, 5.2% in Japan and over 10% in the Baltic countries (IMF, 2010). In some countries, output is expected to continue falling, albeit at a clearly lower rate, also in 2010. Unemployment is rapidly increasing everywhere due to large-scale job destruction. The unemployment rate in the United States and in many European countries is expected to be around 10% in 2010 (IMF, 2010).

Unlike the existing displacement literature, which covers periods in a relatively stable economic environment, the lessons from Finland’s depression period provide a useful point of reference for the size and duration of the earnings losses for workers displaced during the current crisis. Our results suggest that, even in the case of a rapid return to steady economic growth, millions of people losing their jobs during the current downturn will suffer from substantial earnings losses

for several years. Our finding of the large and long-lasting displacement effect at the lower end of the earnings distribution implies that many job losers will be at risk of being outside paid work several years later, possibly withdrawing from the labour market altogether. This is a serious concern, especially for many European governments that have tried to induce people to retire later in order to cope with the financial pressure resulting from an ageing population. Our results underline the importance of supportive measures, such as training and subsidized jobs, to keep job losers employable over the recession period in order to minimize the number of early labour market withdrawals.

From the methodological point of view, our analysis suggests that the mean effect alone can give a rather incomplete picture of the consequences of job loss. For example, a moderate effect of displacement on expected earnings may hide a notable effect that is present only in the left tail of the distribution, as in our 1997 sample.

References

- Angrist, J. D. (2001): Estimation of Limited Dependent Variable Models with Dummy Endogenous Regressors: Simple Strategies for Empirical Practice. *Journal of Business & Economic Statistics* 19, 2–15.
- Appelqvist, J. (2007): Wage and Earnings Losses of Displaced Workers in Finland. VATT Discussion Papers 422.
- Bender, S. – Dustmann, C. – Margolis, D. – Meghir, C. (2002): Worker Displacement in France and Germany. In Kuhn, P. J. (ed.): *Losing Work, Moving On: International Perspectives on Worker Displacement*. W. E. Upjohn Institute for Employment Research, Kalamazoo, 375–470.
- Borland, J. – Gregg, P. – Knight, G. – Wadsworth, J. (2002): They Get Knocked Down: Do They Get Up Again? In Kuhn, P. J. (ed.): *Losing Work, Moving On: International Perspectives on Worker Displacement*. W. E. Upjohn Institute for Employment Research, Kalamazoo, 301–374.
- Burda, M. C. – Mertens, A. (2001): Estimating Wage Losses of Displaced Workers in Germany. *Labour Economics* 8, 15–41.
- Carneiro, A. – Portugal, P. (2006): Earning Losses of Displaced Workers: Evidence from a Matched Employer-Employee Dataset. IZA Discussion Paper No. 2289.
- Couch, K. A. (2001): Earnings Losses and Unemployment of Displaced Workers in Germany. *Industrial and Labor Relations Review* 54, 559–572.

- Couch, K. A. – Placzek, D. W. (2010): Earnings Losses of Displaced Workers Revisited. *The American Economic Review* 100(1), 572–589.
- Doksum, K. (1974): Empirical Probability Plots and Statistical Inference for Nonlinear Models in the Two-Sample Case. *Annals of Statistics* 2, 267–277.
- Eliason, M. – Storrie, D. (2006): Lasting or Latent Scars? Swedish Evidence on the Long-Term Effects of Job Displacement. *Journal of Labor Economics* 24, 831–856.
- Gibbons, R. – Katz, L. F. (1991): Layoffs and Lemons. *Journal of Labor Economics* 9, 351–380.
- Hijzen, A. – Upward, R. – Wright, P. (2006): Using Linked Employer-Employee Data to Estimate the Earnings Costs of Business Closure in the UK. In Bryson, A. – Forth, J. – Barber, C. (eds.): *Making Linked Employer-Employee Data Relevant to Policy*. DTI Occasional Paper No. 4, 134–153.
- Holland, P. (1986): Statistics and Causal Inference (with discussion). *Journal of the American Statistical Association* 81, 945–970.
- Honkapohja, S. – Koskela, E. (1999): The Economic Crisis of the 1990s in Finland: Discussion. *Economic Policy: A European Forum* 14, 401–436.
- Huttunen, K. – Møen, J. – Salvanes, K. J.: How Destructive Is Creative Destruction? The Costs of Worker Displacement. Forthcoming in *Journal of European Economic Association*.
- IMF (2010): World Economic Outlook, April 2010, International Monetary Fund.
- Jacobson, L. S. – LaLonde, R.J. – Sullivan, D.G. (1993): Earnings Losses of Displaced Workers. *American Economic Review* 83, 685–709.
- Jolliffe, D. – Krushelnytskyy, B. – Semykina, A. (2000): Censored Least Absolute Deviations Estimator: CLAD. *Stata Technical Bulletin*, STB–58, 13–16.
- Koenker, R. – Bassett, G. (1978): Regression Quantiles. *Econometrica* 46, 33–50.
- Kletzer, L. G. (1989): Returns to Seniority after Permanent Job Loss. *The American Economic Review* 79, 536–543.
- Kyyrä, T. – Ollikainen, V. (2008): To Search or Not To Search? The Effects of UI Benefit Extension for the Older Unemployed. *Journal of Public Economics* 92, 2048–2070.
- Lehmann, E. L. (1974): *Nonparametrics. Statistical Methods Based on Ranks*. San Francisco: Holden-Day.

- Margolis, D. N. (1999): Part-Year Employment, Slow Reemployment, and Earnings Losses: The Case of Worker Displacement in France. In Haltiwanger, J. – Lane, J. – Spletzer, J. R. – Theeuwes, J. – Troske, K. (eds.): *The Creation and Analysis of Employer-Employee Matched Data. Contributions to Economic Analysis*. Amsterdam: North-Holland, Elsevier Science, 375–416.
- Powell, J. (1986): Censored Regression Quantiles. *Journal of Econometrics* 32, 143–155.
- Ruhm, C.J. (1991): Are Workers Permanently Scarred by Job Displacements? *American Economic Review* 81, 319–324.
- Stevens, A. H. (1997): Persistent Effects of Job Displacement: The Importance of Multiple Job Losses. *Journal of Labor Economics* 15, 165–188.
- Wooldridge, J. M. (2002): *Econometric Analysis of Cross Section and Panel Data*. Cambridge, Massachusetts: The MIT Press.

4. Institutional rules, labour demand and retirement through disability programme participation⁴⁷

Abstract

We use matched employer-employee data from Finland to model transitions out of work into sick leave and disability retirement. To identify the role of institutional factors we exploit reforms that changed medical requirements for disability pension eligibility and experience-rated employer contributions. We find that transitions to sick leave and disability pension benefits are relatively rare in growing establishments, but rather common in establishments with a high degree of excess worker turnover. We also show that transitions to disability retirement depend on the stringency of medical screening and the degree of experience-rating applied to the employer.

Key words: Disability pension, sick leave, experience rating

JEL classification numbers: J14, J23, J26

4.1 Introduction

The disability benefit scheme is one of the largest social security programmes in many countries, and therefore is of particular interest. In Finland, disability is the most common reason for early retirement, and disability expenditure accounted for some 3.5% of GDP in 2003, which was the third highest share in the EU after Sweden and Denmark (Börsch-Supan, 2007). Disability enrolment rates of older employees vary strikingly across the European countries and the US. These cross-country differences cannot be explained by demographic or health-related factors, but are attributable to institutional differences in the disability schemes (Börsch-Supan, 2007). During the past two or three decades, many countries have also experienced an expansion of disability benefit enrolment. This is a serious concern given the common goal to induce people to retire later. The

⁴⁷ Ossi Korkeamäki and Tomi Kyyrä.

Financial support from the Finnish Centre for Pensions is gratefully acknowledged. We are grateful to Mikko Kautto, Raija Gould and Juha Rantala for their helpful comments and to Tuuli Ylinen and Eija Mustonen for their research assistance. The comments of two anonymous referees and a co-editor substantially improved the paper. This paper is forthcoming in the *Journal of Population Economics*. An earlier version was published in the Finnish Centre for Pensions Working Papers series (No 2009:5). The final publication is available at www.springerlink.com (DOI: 10.1007/s00148-010-0330-z).

widespread use of disability benefits as an early retirement instrument has been argued to be a particularly serious problem in Finland (e.g. OECD, 2008).

Disability benefits are designed to provide insurance for employees' labour income against the risk of becoming disabled and incapable of regular work. In practice, it may be difficult to identify employees who are truly disabled, which suggests the possibility that disability benefits can distort labour supply and demand in some cases. Autor and Duggan (2003, 2006), for example, argue that the rapid growth in disability benefit rolls in the US cannot be explained by changes in health, but is driven by a combination of labour demand conditions and changes in the disability scheme itself (in terms of generosity, coverage and screening intensity). Other authors have also found evidence of the importance of the generosity of disability benefits, the stringency of medical screening and the economic environment when explaining participation in disability programmes (e.g. Gruber, 2000, Black *et al.*, 2002, and Campolieti, 2004). A majority of this literature has been motivated by a simple labour supply framework, in which an employee chooses whether to apply for disability benefits, while the employer has no role at all. Surprisingly little effort has been made to study the labour demand side (some exceptions are Hassink *et al.*, 1997, and Koning, 2009).

When job cuts are necessary, firms often offload their oldest employees first. If the health requirements for disability benefit eligibility are weak, early retirement via the disability scheme can be a useful strategy in effective downsizing, providing a way to reduce the workforce in a "soft" way. On the other hand, some firms can also target dismissals at those employees with a high risk of disability. In doing so, the employer may avoid disability costs arising from the experience-rated contributions of disability pension benefits. Encouraging disability retirement can also be an attractive strategy for an employer who wants to change the composition of the workforce at a time of stable or growing employment when dismissals are difficult to justify.

This study aims to shed some light on the relationships between labour demand, institutional factors and early retirement through disability programmes. We consider the importance of the labour demand side by examining the relationship between the establishment's growth and restructuring rates and disability entries by its employees. In addition, we assess the effectiveness of two policy instruments: the strictness of medical requirements for disability pension eligibility and the experience-rating of disability expenditure. The first determines the ease of access to disability pension benefits, whereas the latter places part of the costs of early retirement on the employer.

Using matched employer-employee data from Finland, we model transitions out of work to sick leave and disability retirement. To identify the role of institutional factors we exploit a law change that made the medical requirements for disability pension eligibility tougher for a certain group, as well as changes in

partially experience-rated employer contributions. We show that transitions to sick leave and disability pension benefits are relatively rare in growing establishments, but rather common in establishments with a high degree of excess worker turnover. We find no evidence of employers actively encouraging disability retirement as a way of adjusting their workforce when downsizing. Finally, we show that the transition rate to disability retirement depends on the stringency of medical screening and the degree of experience-rating applied to the employer.

The paper proceeds as follows. In the next section, we give a short overview of the existing literature. Section 4.3 describes the Finnish social security system for sickness benefits and disability pensions. We discuss our data and report descriptive statistics in section 4.4. The results of our econometric analyses are reported and discussed in section 4.5, which is followed by a concluding section.

4.2 Related literature

Disability benefits are typically determined as a function of past earnings, which are likely to be correlated with the employee's preferences for work. The resulting endogeneity problem has hampered attempts to quantify the impact of disability benefits on labour supply. Gruber (2000) and Campolieti (2004) overcome the endogeneity problem by exploiting policy changes in the Canadian disability benefit scheme that had differential effects on people living in different parts of the country. While Gruber estimates that the elasticity of nonparticipation with respect to disability insurance benefits is between 0.28–0.36, Campolieti finds no statistically significant relationship. In the late 1980s, which is the period analysed by Gruber, non-medical factors related to the availability of suitable jobs in the region and personal skills were taken into consideration when determining benefit eligibility. Campolieti considers an earlier reform that took place in the early 1970s when the eligibility requirements and the stringency of medical screening were tougher. This led Campolieti to argue that the generosity of the disability benefits may not distort working decisions when it is difficult to qualify for such benefits because of a strict screening process.

Using aggregate data for the US, Black *et al.* (2002) and Autor and Duggan (2003, 2006) find evidence of the relationship between disability participation and business cycle conditions. Black *et al.* use data from the coal boom and bust in the 1970s and 1980s, which affected only a few coal-producing counties, to construct instrumental variables for local labour market shocks. According to their county-level analysis, participation in disability programmes falls during economic upturns, and this relationship is much stronger for permanent than for

transitory economic shocks. Autor and Duggan discuss a dramatic expansion of disability insurance enrolment during the past two decades in the US.⁴⁸ They argue that this growth cannot be explained by a true increase in the incidence of disabling illness. Instead, they claim that the reduced stringency of the screening for disability benefits after 1984, an increase in the earnings replacement rate, and an increase in female labour force participation have played important roles. Because of the liberalization of disability benefits in 1984, the disability application rates were found to become more responsive to adverse labour demand shocks. This result supports Campolieti's (2004) interpretation of Canadian evidence. Autor and Duggan (2003) estimate that the unemployment rate of workers aged 25–64 in 1998 was a half percentage point lower than it would have been otherwise. They argue that the US disability system has begun to “function much like a long-term unemployment insurance programme for the unemployable” (Autor and Duggan 2006, p. 74).

Unlike the US studies based on aggregate data, Vahtera *et al.* (2005) and Rege *et al.* (2009) analyse individual-specific disability risks. Using matched employer-employee data for Norway, Rege *et al.* explain the likelihood of being on a disability pension with dummy variables indicating various degrees of plant downsizing during the past six years. They find a substantial increase in the transition rate to disability retirement following plant downsizing or plant closure.⁴⁹ Whereas Rege *et al.* do not make a distinction between those who kept their jobs and those who lost their jobs in plant downsizing, Vahtera *et al.* consider a risk of disability retirement among Finnish municipal employees who kept their jobs after the reduction of personnel in their organisation. These employment reductions were carried out between 1991 and 1993, during a period of severe recession in Finland. They find an almost twofold risk of being granted a permanent disability pension in the next five years after a major downsizing (more than 18% reduction in the personnel) than after no downsizing (less than 8% reduction). Thus, not only employees who lose their jobs, but also those who keep their jobs after the employment reduction are subject to an increased disability risk.

Börsch-Supan (2007) points out that disability expenditures and enrolment rates vary notably across different countries. In Europe disability expenditures are highest in Finland, Sweden and Denmark.⁵⁰ Börsch-Supan analyses the cross-

⁴⁸ See McVicar (2008) for discussion about the growth in disability benefit rolls in the UK.

⁴⁹ They also find an increase in mortality rates among workers whose plants downsized.

⁵⁰ Disability schemes represent only a part of the social security system. How people who are unable, or unwilling, to work are allocated between sickness, unemployment, disability, and early retirement schemes depends on relative compensation levels and eligibility criteria, which vary from country to country. Hence, a low disability enrolment rate may be associated with a high rate of sickness absence,

country differences in the disability enrolment rates of people aged 50 to 65 years, using harmonized survey data for 12 European countries and similar survey data for the UK and US. He finds very little explanatory power for demographic and health-related differences across the countries. By contrast, three quarters of the cross-country variation was explained by the institutional variables that describe the generosity and the ease of access of the disability insurance. The most influential institutional factor turned out to be the strictness by which vocational considerations are applied when determining eligibility.

The studies discussed above do not pay much attention to the employer's role. Hutchens (1999) develops a theoretical framework that helps to understand why employers may be actively involved in early retirement decisions. He introduces an implicit contract model of a firm that uses early retirement benefits, provided by the government, as a form of unemployment insurance. Within this framework, the public early retirement benefits effectively subsidize workforce reductions. Therefore, the firm responds to slack demand by Encouraging early retirement, which leads to an inefficiently high level of early retirement. Hutchens also discusses two alternative policies: actuarial adjustments and experience rating. An actuarial adjustment places costs on early retirees by reducing their future benefits compared to the case where retirement occurs at a later day, whereas experience rating places costs on firms by directing part of the early retirement expenditure to the former employer. While both of these policies can be used to reduce the implicit subsidy, and thereby restore early retirement to the efficient level, their implementation is subject to some practical drawbacks. Namely, an effective early retirement scheme should vary with individual characteristics, like wages and survival probabilities. Since the real-world scheme cannot account for all individual heterogeneity, the implicit subsidy will exist at least for some groups even if the scheme eliminates the subsidy "on average".

Theoretical insights of Hutchens (1999) are supported by empirical findings of two studies from the Netherlands. Using data on the dismissal and disability rates of Dutch firms, Hassink *et al.* (1997) examine to what extent separations into disability are used as an alternative to dismissals. They estimate that about one-tenth of the observed inflow into disability were effectively dismissals. The data used by Hassink *et al.* covered the years 1988 and 1990 when the experience-rating of disability benefits was not yet introduced in the Netherlands. In 1998 the employer's annual disability insurance contribution rate was tied to the amount of disability benefits received by its former employees during a past 5-year period (beginning 7 and ending 2 years prior to the year in question). As a

long-term unemployment or voluntary early retirement. These kinds of spillovers should be kept in mind when interpreting the results from cross-country comparisons.

result, the employers became partly liable for the costs of the first five years of disability benefits. Koning (2009) compares inflow rates to disability benefits between employers that experienced a change in the contribution rate in 2001 (triggered by a decline or increase in the disability inflow in 1999 compared with an earlier period) and those with no change in the contribution rate. He finds that in the firms that experienced a (positive or negative) change in the contribution rate the disability inflow rate decreased during the next two years compared with the firms with no change in the contribution rate. Koning interprets this as evidence that employers were not completely aware of the experience rating scheme, and hence the change in the contribution rate acted as a “wake-up call” to pay attention to experience rating. This in turn induced the employers to increase preventive actions, reducing disability events in the subsequent years.

It should be stressed that Koning’s data only covered three post-reform years, and thereby there were no exogenous changes in the experience rating scheme during the observation period. Instead, all the observed changes in the contribution rates were driven by changes in firms’ own disability history. While Koning’s findings indicate some information imperfections (at least a few years after the introduction of the experience-rating system) and imply that experience rating does matter, his results do not describe the causal effects of having a given degree, or a particular type, of an experience-rating system compared with the counterfactual case of having some other scheme.

In sum, we can draw the following lessons from the existing literature: 1) the generosity of disability compensation and negative demand shocks increase the entry rates to disability benefit schemes, 2) the strength of this relationship depends on the stringency of medical screening, and 3) the experience-rating of disability benefit costs can be used to reduce the moral hazard problem. Our study complements this literature in a number of ways. First, in addition to studying transitions from work to disability retirement, we also consider transitions from work to sick leave and from sick leave to disability retirement. In this way, we can differentiate between factors affecting the incidence of sickness or injury (*ex ante* effects) and those affecting the intensity or success of medical and occupational rehabilitation (*ex post* effects). This distinction helps us to detect the point in the disability pension track when certain policy instruments are effective. Second, when analysing the role of labour demand, we pay particular attention to excess worker turnover, which describes a degree of restructuring for a given employment change. This helps us to show that disability benefits may be used to adjust the structure of the workforce at the times when the employment level is stable. Third, we take advantage of two policy reforms to identify the causal effects of the experience-rating of disability benefits. Our results indicate that experience rating reduces the incidence of disability and sickness. Finally, by accounting for a firm’s financial position, we also show that experience rating has a heterogeneous effect, being less effective

for those employers that can easily bear their share of the disability pension costs.

Given that there is hardly any evidence on the experience-rating of disability benefits, our analysis of the effects of experience-rated contributions on the disability entry rate is the main contribution of the paper. Koning (2009) is an exception, but we extend his work in several ways. Most importantly, our estimates for the impact of experience-rating can be given a causal interpretation. The lack of prior evidence is partly due to the fact that disability benefit expenditures are subject to experience rating only in a few countries. Still, the topic should be of considerable interest as many countries suffer from high and still growing rates of disability programme participation, and experience rating is one potentially effective policy instrument.

4.3 Institutional framework of Finland

The Finnish social security system has been subject to continuous changes over time. Below we describe the features of the system that were in force from the early 1990s until 2004, which is the time period covered by our empirical analysis.

4.3.1 Sickness and disability benefits

An employee who is unable to perform his job due to illness or injury is entitled to compensation for income losses. The applicant needs a statement by a doctor or hospital certifying that he is not capable of work. For the first ten working days the applicant is fully compensated by the employer, after which he can claim a *sickness benefit* from the Social Insurance Institution (KELA). Depending on the collective labour agreement, many employers continue to pay wages or salary after the mandatory waiting period of ten working days, in which case the allowance is paid to the employer. As a result, the time out of work until receipt of a sickness benefit directly from the Social Insurance Institution is typically one to three months. The sickness benefit is determined by the past taxable earnings, and it can be received for a maximum of about one year (300 working days, Saturdays included). Depending on illness or disability, the applicant's rehabilitation needs and possibilities are assessed in a more extensive medical examination during the sickness benefit period. In case of a prolonged illness or permanent disability, the employee can apply for a disability pension.

An *ordinary disability (OD) pension* is payable to individuals aged 16 to 64 whose working capacity has significantly decreased. A full benefit is conditional on the working capacity loss of at least 60% and a partial benefit for a loss of 40 to 59%. When determining eligibility, an individual's capability to support herself by regular work, age, education, occupation, and place of residence are taken into account along with the medical assessment. The OD pension can be

granted either indefinitely (if return to work is not likely) or for a specific period. In the latter case, the OD pension is also referred to as a rehabilitation subsidy or a cash rehabilitation benefit, and its receipt is conditional on a rehabilitation plan. An OD pension may be discontinued if the working capacity of the recipient improves, but this rarely happens among older recipients (e.g. OECD, 2008, p. 116). There is no automatic retesting of the disability status except for new periods of the rehabilitation subsidy.

An *individual early retirement (IER) pension* is another disability pension, which is available for employees who have a long working career and who are unable to continue in their current job because of deteriorated health. Compared with the OD pension, eligibility for the IER pension is subject to less strict medical criteria. The minimum degree of working incapacity is not defined and occupational factors like the length of service and working conditions carry greater weight. It suffices that working capacity has reduced to such an extent that the person cannot continue in her present job or occupation, so that other working possibilities are not considered. Unlike the OD pension however, the IER pension is payable only to employees above a certain age threshold. In the private sector, there was a uniform age threshold of 55 until 1994 when it was raised by three years to 58 for people born in 1940 or later. In 2000 the age threshold was raised further by two years to 60 for those born in 1944 or later. In 2004 the IER scheme was abolished entirely from these same cohorts. At the same time the medical criteria for OD pension eligibility were somewhat relaxed for people aged 60 and over.

Saurama (2004) provides some survey evidence that bad health is not the only reason for entering into a disability pension in Finland.⁵¹ As expected, disability pensioners reported bad health as one of the main reasons for retirement, but many of them said that straining work played an important role as well. In particular, 62% of OD pensioners and 74% of IER pensioners had felt that their job had become too exhausting or they could not handle their job any more. A notable fraction of the respondents had also felt pressure to retire from the management or colleagues: 14% of OD pensioners and 24% of IER pensioners were partly forced out of their job. Hence, difficulties in performing job tasks and the pressure from the workplace are important factors affecting disability pension entry.

⁵¹ The target population of the survey was all people who received early retirement benefits at the end of 1998. The early retirees were asked for their reasons of retiring. The response rate was quite low – only 51.3% – and young disability pensioners were under-represented. The numbers referred to in the text were taken from table 12 in Saurama (2004, p. 132).

The purpose of the gradual abolition of the IER pension was to reduce the disability enrolment rates at higher ages. This reform made the medical requirements for disability pension eligibility tougher for the later cohorts who have been able to apply only for the OD pension, and thereby should reduce the flow into disability pension benefits among those who are not truly disabled. Because only those employees born after 1939 (1994 reform) or after 1943 (2000 reform) were affected, these reforms provide us with a quasi-experimental setting for studying the importance of the stringency of medical screening.

4.3.2 Experience rating of disability pension benefits

A particular feature of the Finnish disability scheme is that employers are partially liable for the disability pension costs of their former employees via experience-rated employer contributions. Experience rating is not applied to firms with fewer than 50 (300 until 1995) employees, which pay a fixed tariff rate for each employee. The larger firms are partially covered by experience-rated contributions and partially by fixed (age-dependent) tariff rates. The employer subject to experience rating must pay its share of the present value of disability pension costs at the time when a disability pension is awarded to its former employee. Given that disability pension costs can accumulate over several years until the person reaches age 65 and transfers to an old-age pension, the disability event can become very costly for the former employer in the case of a large firm.

To be more specific, consider an employee i of firm j who is awarded a disability pension in year t . Therefore, his employer has to make a lump-sum contribution equal to

$$C_{j,i} = \alpha(\text{size}_{j,t-1})\gamma(\text{age}_{i,t})b_i, \quad (20)$$

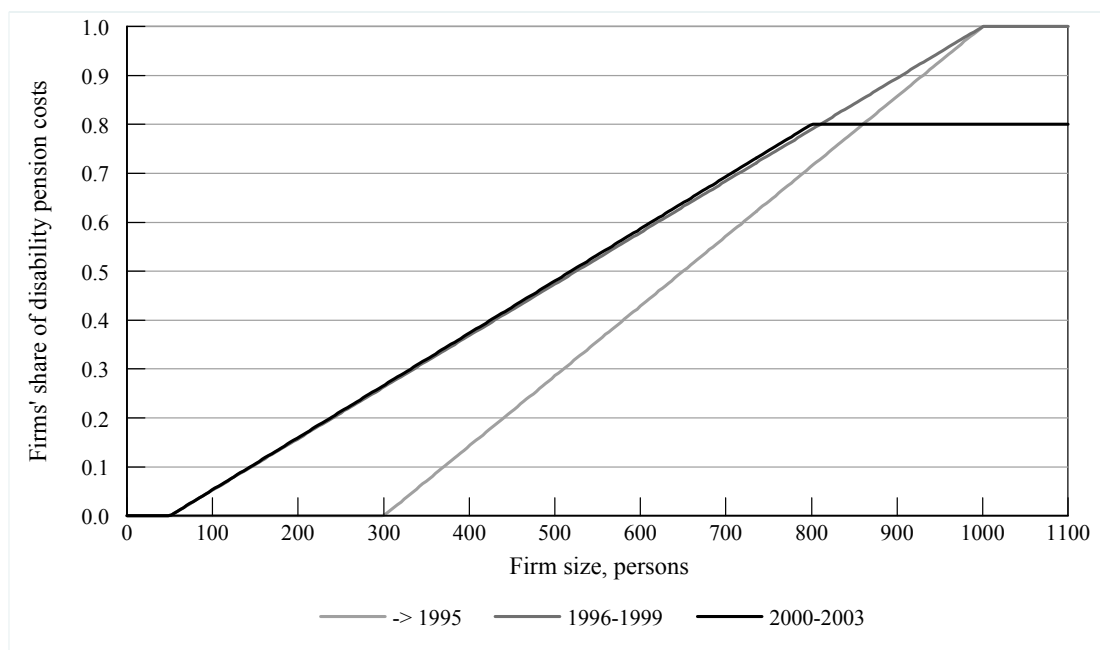
where b_i is an annual disability pension benefit, g is the present-value multiplier and α is the degree of experience-rating applied to firm j . The product of pension benefit b_i and multiplier g serves as an estimate of the present value of expected disability pension benefits up to the age when the entitlement to an old-age pension begins.⁵² The multiplier g is a decreasing function of worker's age at the time of disability retirement, ranging from 9.66 at age 50 to 2.31 at age 62 for the groups analysed in this study.

⁵² We obtained the values of γ from a pension institution. All pension institutions use the same values. The multiplier is determined by the average duration of disability pension receipt of persons who are granted a disability pension at a given age. It also accounts for the (average) probability that the person returns to work and the (average) survival probability until age 65.

The realized marginal cost of a disability retirement entry depends crucially on the degree of experience-rating a , which is determined by the size of the firm's workforce in year $t-1$. Until 1995, a was 0 for all firms with fewer than 300 employees. In these firms the pension contributions were independent of the retirement events of their employees and, consequently, the marginal cost of disability retirement was 0. For the larger firms, a increased as a linear function of firm size from zero (300 employees) to one (1000 employees); see Figure 4.1. In other words, the largest firms with at least 1000 employees were fully liable for the expected disability pension costs of their former employees.

During the time period under investigation, the experience rating scheme changed twice. In 1996 the experience rating scheme was first extended to cover also firms with 50 to 299 employees. A smaller reform took place in 2000 when the maximum liability share was limited to 0.8. As a result of the two reforms, the degree of experience-rating varies across firms of a given size over time, which can be seen in Figure 4.1. By exploiting this variation for identification, we can distinguish the effect of experience-rating from the firm size effect.

Figure 4.1 The degree of experience-rating (a) as a function of firm size in different periods



To highlight the size of disability costs for the employer, consider a worker who is awarded a disability pension at age 55 $g \approx 7.22$. Assuming the pension benefit equals 55% of the past annual salary, which is true on average, the maximum disability cost for the former employer would be as much as four times the annual salary (i.e. $\alpha = 1$ for a firm with over 1000 employees before 2000). In the case of a firm with 400 employees, the disability cost would correspond either to 7 months' salary (pre-1996 rules) or 18 months' salary (rules from 1996

onwards). Hence, it is evident that the disability costs for the former employers can be quite large, and that these costs changed substantially due to the two reforms.⁵³

The aim of experience-rating is to minimize the employer's moral hazard problem by placing costs on those firms whose employees enter the disability pension schemes. When an employee applies for a disability pension, the employer has no direct control over the decision made by the pension institution. Nevertheless, the employer has the means to influence indirectly the flow into sick leave and the likelihood that recipients of sickness benefits will return to work rather than retire via a disability pension scheme. If effective, experience rating should induce the employer to take preventive measures to minimize the flow into sick leave (the *ex ante* effect), as well as to put effort into getting its employees back to work from sick leave (the *ex post* effect). The preventive action may involve reallocation of the workload to minimize stress-related illness and arrangements that reduce accidents at the workplace. When helping people come back to work from sick leave, occupational rehabilitation and job modifications that allow the switching of jobs within the firm are crucial for those who cannot perform their old tasks despite medical rehabilitation.

4.4 Data and descriptive evidence

Our data set was drawn from the records of the Finnish Longitudinal Employer-Employee Database (FLEED). Employee information in the database is obtained by merging information from over 20 administrative registers with unique personal identity numbers. The database covers effectively everyone with a permanent residence in Finland. Along with standard socio-demographic background variables, the database includes detailed information on annual income (from the tax authorities), job spells (from the pension institutes), unemployment spells and participation in labour market programmes (from the employment offices). For people who are employed in the last week of a given year, the ES database also includes the unique identification code of the firm and establishment. This allows us to identify individuals who are working for the same employer and provides a link to firm records. Thus we are able to measure labour turnover and employment changes at establishment and firm levels.

The principal source of firm records in the FLEED is the Financial Statements Statistics (FSS), which is an annual survey conducted by Statistics Finland. The

⁵³ It is worth noting that the Finnish experience-rating system differs from the Dutch one studied by Koning (2009) at least in the following ways: 1) the employer's liability is not limited to the first five years of disability benefit costs, 2) the disability event causes a lump-sum payment, having no effect on the pension contributions thereafter, and 3) the degree of experience-rating varies much more across firms of different size.

survey contains corporate income statement and balance sheet data on firms in manufacturing, construction, retail and wholesale trade, business services, hotel and restaurant services, and transportation. This data is available with time consistent variable definitions for the period 1986–2005. All firms above a certain size threshold, which varies between the sectors and over time, have been included in the survey. Until 1996 also a sample of smaller firms was included in the survey, but since then Statistics Finland has collected information on the small firms only from the administrative registers. So, the survey data for the later years have been complemented by adding firm records from the Business Tax Register with more limited information content but covering all firms in the private sector. The combined survey-register data should be dynamically representative over all firms in each year, although some small firms are missing from the first six years of our observation period.

Some key variables in the employee data contain information on sickness benefits, paid by KELA during the year, and on the types of pension benefits received at the last week of the year. These are used to detect transitions into sick leave and disability retirement. As discussed above, employees on sick leave are fully compensated by their employers for the first ten days to three months, depending on the collective labour agreement under which they are employed. Hence, receipt of sickness benefits directly from KELA indicates a prolonged illness.


We can distinguish between OD and IER pension benefits, but we do not know the compensation level (partial or full benefit), or whether a disability pension was granted indefinitely or for a specific period. We classify an employee as being on a disability pension if he or she received either OD or IER pension benefits at the last week of the year.

4.4.1 Incidence of disability retirement

We begin by considering the extent of the disability problem in Finland. While the official retirement age was 65 until 2005⁵⁴, the effective retirement age – the average age of new pensioners – has been around 60 due to early retirement schemes, of which the disability schemes are the most important ones. In 2007 a roughly equal number of new pensioners were granted a disability pension and an old-age pension, but a few years earlier the disability pension was the most common pathway to retirement. This explains why disability expenditures are so high in Finland compared with other industrialized countries.

⁵⁴ The ordinary old-age pension was available for people older than 61 but only those entering at age 65 received the full benefits. Since 2005, employees have been able to choose freely at which age between 63 and 68 they begin to collect the old-age pension benefits.

Figure 4.2 Population share of disability pensioners by birth year



Note: Age at the end of the year used as a marker for men. Numbers for 1995 are not shown due to data errors. Source: Authors' calculations from the FLEED.

A high incidence of disability is illustrated in Figure 4.2 where the disability enrolment rates (including both OD and IER pension recipients) are shown as a function of age for four birth cohorts by sex.⁵⁵ Within all birth cohorts, women have lower enrolment rates at each age.

Compared with the later cohorts, employees born in 1940 are more likely to be disabled at all ages. Close to 30% of this cohort were on a disability pension at age 60, which is a strikingly large figure. There are no notable differences in disability rates at a given age between the 1945, 1950 and 1955 cohorts. To some extent the lower disability rates for these cohorts may be related to their ineligibility for the IER pension scheme.

4.4.2 Outcome variables for analysis of transitions

We distinguish the likelihood of entering sick leave, which often means only temporary absence from work, from the likelihood of being granted a disability pension, which almost surely means a permanent withdrawal from the labour

⁵⁵ Statistics Finland changed its procedure of merging register data on pension benefits in 1995. This led to an unexplained (small) drop in the number of disability pension recipients in the FLEED for that year, reflecting some technical problems. For that reason we chose to break the time series in 1995.

market. The determinants of these events can differ and be affected by experience rating in different ways. Hence, we shall model transitions from work to a disability pension, from work to sick leave, and from sick leave to a disability pension. This approach ignores possible spill over effects toward other exit destinations, most notably into unemployment, which should be kept in mind when interpreting our results.

In all of our models the risk set in year t includes employees who (i) were 50 to 62 years old at the end of year $t - 1$, (ii) held a job at the end of year $t - 1$ in a private-sector firm with at least ten employees, (iii) had been working, without receiving any pension benefits, at least for three consecutive years (i.e. from the beginning of year $t - 3$ until the end of year $t - 1$), and (iv) did not receive sickness benefits during year $t - 1$.⁵⁶ We exclude the younger and older employees because their transitions into disability retirement are very rare. Moreover, for younger employees it is probably very difficult to be granted a disability pension without serious injury or illness, whereas the older employees can retire via other early retirement schemes, suggesting that the misuse of the disability pension schemes is not a serious issue for these groups. We also exclude employees from firms with less than ten employees, as the data on very small firms is noisy.

Transitions to disability pension. Because receipt of a disability pension typically follows a sick leave and/or rehabilitation measures, there is a gap between the job withdrawal, which is of our primary interest, and actual entry into disability pension benefits. To detect the year when the process towards disability retirement started, we follow each person at risk in year t for the next three years (two years from 2002). The employee is classified as becoming disabled in year t if her working career was interrupted during that year and she was granted a disability pension by the end of year $t + 2$, without being unemployed or employed in the another firm in meantime.⁵⁷ In other words, we are interested in transitions from a given workplace to disability retirement, but allow for periods of sick leave and rehabilitation between these two events. The majority of disability pensions following job withdrawal in year t are granted during year t (46%) or $t + 1$ (48%), while the number of entries into disability retirement drops sharply in year $t + 2$. A few pensions are also granted at the later

⁵⁶ The cost of disability pension is borne by the former employer (according to the experience-rating rules concerning firm size) only when the employment relationship has lasted for a minimum of three years. We also include workers who changed their jobs within 3-year period but control for job tenure in the probability models. Excluding these workers from the analysis does not notably affect our results.

⁵⁷ Receipt of a disability pension is not observed for some people in 1995 due to the change in the procedure of merging the underlying register data. Using the 3-year moving window for transitions to disability pension also minimizes this problem.


periods but we think that assigning them to the employer at the end of year $t - 1$ is too unreliable and hence we discard those cases.

Transitions to sick leave. Receipt of an OD pension typically follows a one-year period on sickness benefits, and therefore almost all OD pensioners have been on sick leave before retiring. However, some employees have retired directly via the IER scheme without being on sickness benefits first. Since deteriorated health is a prerequisite for receipt of an IER pension as well (and since short spells of sickness benefits are not observed in the data), also these employees are classified as entering sick leave in the year when the pension was granted. Thus, an employee at risk in year t moves into sick leave during that year if he started to collect sickness benefits or was granted a disability pension during year t even without actually receiving sickness benefits. About a quarter of workers who were awarded a disability pension did not receive sickness benefits. However, apart from the effect of IER pension eligibility, our estimates appear to be fairly robust with respect to the treatment of this group.

Transitions from sick leave to disability pension. When modelling the likelihood of being granted a disability pension conditional on being on sickness benefits, the risk set in year t includes only those employees who started to collect sickness benefits during year t , including also those who were granted a disability pension without a period of sick leave.

Illustration. Four possible labour histories are shown in Figure 4.3. Each of these persons is at risk of making a transition to disability retirement and to sick leave in year t . Employees A, B and C withdrew from work in year t and were granted a disability pension by the end of year $t + 2$. According to our definition, they made a transition from work to disability retirement in year t ($D_t = 1$). Since employees A and D received sickness benefits and employee C started to collect disability pension benefits in year t , they are also classified as entering sick leave in that year ($S_t = 1$). When modelling transitions from sick leave to disability retirement, employees A, C and D would be included in the risk set in year t , but not employee B who was out of work for an unknown reason.

Figure 4.3 *Examples of labour market histories for persons at risk in year t .*



4.4.3 Sample design for modelling transition rates

The transition probabilities (or rates) defined above are closely interrelated but not perfectly. Namely, it holds that

$$\Pr(D_t = 1) \geq \Pr(S_t = 1)\Pr(D_t = 1 | S_t = 1), \quad t = 1991, 1992, \dots, 2002,$$

where $\Pr(D_t = 1)$ denotes the likelihood of starting disability retirement (by the end of year $t + 2$) due to job withdrawal in year t , $\Pr(S_t = 1)$ is the likelihood of entering sick leave in year t and $\Pr(D_t = 1 | S_t = 1)$ is the likelihood that a sick leave starting in year t eventually lead to disability retirement. It should be stressed that employees whose working career was interrupted without receipt of sickness benefits in year t and who were granted a disability pension in year $t + 1$ or $t + 2$ contribute to the event on the left-hand side of the equation but not to the events on the right-hand side (e.g. person B in Figure 4.3). As a result, the product of the two probabilities on the right-hand side is typically less than the overall disability probability on the left-hand side. Nevertheless, the product of $\Pr(S_t = 1)$ and $\Pr(D_t = 1 | S_t = 1)$ gives a useful decomposition for $\Pr(D_t = 1)$ although its approximate nature should be kept in mind when interpreting the results. By comparing the determinants of $\Pr(S_t = 1)$ and $\Pr(D_t = 1 | S_t = 1)$, we aim at making a distinction between the *ex ante* and *ex post* effects of the covariates of interest.

Our estimation samples include all employees who are at risk of making a transition of interest between the years 1991 and 2002. The sample period ends in 2002, because when detecting the timing of job withdrawals, we have to be able to follow employees at risk at least for the next two or three years. The first period is 1991 because the FLEED employee data are available from 1988 onwards and we need employee records from the past three years to construct the worker flow variables. Our samples are rotating panels where in each year new employees enter the risk set while some old ones leave it.

4.4.4 Raw transition rates

Table 4.1 reports the sizes of the risk groups and the raw transition rates by age. All transition rates increase almost uniformly with age. The likelihood of a transition to sick leave becomes five times larger while the likelihood of being granted a disability pension grows eightfold from age 50 to 62. This is explained by the fact that employees are less likely to return to work from sick leave at older ages, as seen in the last column. The transition rate from sick leave to disability retirement is strikingly high at all ages; one-half of sickness benefit recipients aged 50, and four in five above age 58, do not return to work but end up in disability retirement.

Table 4.1 Transition rates and the size of risk groups by age

Age	N	$\Pr(D = 1)$	$\Pr(S = 1)$	R	$\Pr(D = 1 S = 1)$
50	126,677	0.007	0.012	1,517	0.483
51	120,358	0.008	0.012	1,496	0.499
52	115,200	0.010	0.015	1,673	0.546
53	107,084	0.013	0.017	1,805	0.576
54	97,444	0.018	0.023	2,227	0.599
55	80,917	0.021	0.026	2,119	0.640
56	66,223	0.025	0.027	1,821	0.661
57	52,680	0.033	0.034	1,812	0.728
58	40,484	0.038	0.036	1,463	0.759
59	30,803	0.051	0.045	1,382	0.815
60	22,044	0.055	0.047	1,040	0.845
61	16,456	0.061	0.053	869	0.826
62	10,976	0.059	0.058	634	0.841
All	887,346	0.019	0.022	19,858	0.655

Notes: N is the size of the risk group for transitions to disability retirement and sick leave. R is the size of the risk group for transitions from sick leave to disability retirement. $\Pr(D = 1)$ is the probability of a transition to disability retirement. $\Pr(S = 1)$ is the probability of a transition to sick leave. $\Pr(D = 1 | S = 1)$ is the probability of a transition to disability retirement conditional on being on sick leave.

It is noteworthy that the size of the group at risk of entering sick leave or becoming disabled declines rapidly with age. While the total number of employees at risk is about 127,000 at age 50, the group at risk halves by age 56 and includes only 11,000 employees at age 62. This declining age pattern reflects the fact that people tend to withdraw from employment quite early. In particular, it is rather common for private-sector employees to end up in long-term unemployment at older ages because employees above a certain age threshold at the time of unemployment entry can collect earnings-related unemployment insurance benefits until retirement at age 60 via the unemployment pension scheme (see e.g. Kyrrä and Wilke, 2007).

Because of the experience-rating system, variation in the transition rates across firms of different sizes is of particular interest. In Table 4.2, employees are classified into four groups according to the size of their employer at the end of year $t - 1$. As seen in the last column, there are no systematic differences in the transition rates to sick leave or disability retirement in firms with 1000 or fewer employees. However, in the largest firms the transition rate to both sick leave and disability retirement is notably higher than in the smaller firms. The experience-rating of disability pension expenditures, which devotes higher cost shares for larger firms, would have implied the opposite. Moreover, the transition rate from sick leave to disability retirement does not vary with firm size, giving no support for the *ex post* effect of experience-rating.

Table 4.2 Transition rates by firm size and establishment growth

Firm size	Establishment's growth category						All
	Past period from $t - 4$ to $t - 1$			Current period from $t - 1$ to $t + 2$			
	Contracting	Stable	Expanding	Contracting	Stable	Expanding	
From work to disability retirement, $\Pr(D = 1)$							
5–50	0.016	0.017	0.016	0.016	0.016	0.016	0.016
51–300	0.020	0.018	0.017	0.019	0.019	0.016	0.018
301–1000	0.019	0.018	0.016	0.019	0.018	0.018	0.018
> 1000	0.023	0.022	0.022	0.023	0.022	0.023	0.022
All	0.020	0.020	0.018	0.020	0.019	0.019	0.019
From work to sick leave, $\Pr(S = 1)$							
5–50	0.019	0.020	0.019	0.019	0.020	0.019	0.020
51–300	0.021	0.021	0.019	0.022	0.021	0.018	0.021
301–1000	0.022	0.021	0.019	0.023	0.020	0.020	0.021
> 1000	0.028	0.026	0.025	0.027	0.025	0.027	0.026
All	0.023	0.023	0.021	0.023	0.022	0.022	0.022
From sick leave to disability retirement, $\Pr(D = 1 S = 1)$							
5–50	0.594	0.615	0.574	0.606	0.605	0.574	0.598
51–300	0.703	0.659	0.645	0.676	0.669	0.640	0.666
301–1000	0.675	0.667	0.655	0.655	0.666	0.683	0.666
> 1000	0.656	0.672	0.643	0.671	0.667	0.632	0.662
All	0.661	0.661	0.634	0.659	0.660	0.631	0.655

Establishments were divided into contracting, stable and growing ones according to 3-year growth rates from $t - 4$ to $t - 1$ and from $t - 1$ to $t + 2$. In contracting establishments the growth rate was below -0.1 , in stable ones between $[-0.1, 0.15)$ and in growing ones greater or equal to 0.15 .

There are reasons why disability and sick leave entries may be associated with current or past employment changes at the workplace. The past employment reductions can lead to added stress at the workplace such as reduced control over one's chores and increased workload and job insecurity. This could cause health problems for those who kept their jobs (e.g. Vahtera *et al.*, 2005). Encouraging disability retirement of older employees can be a helpful strategy in downsizing and restructuring endeavours of a firm, suggesting that there might be a relationship between the current growth rate and disability incidence.⁵⁸ Since the

⁵⁸ A recent study by Gielen and Van Ours (2006) analyses age differences in job reallocation and labour mobility using matched worker-firm data for the Netherlands. They find that firms adjust their workforce mainly via entry for young and prime-age workers, but via separations for older workers. Furthermore, employment of old workers is found to be more responsive to firm-specific employment changes.

difficulties of laying off older employees probably vary with business conditions, the disability entry rates can also vary between downsizing and expansion periods.

To address these questions we examine the relationship between the transition probabilities and the past and current growth rates of employment. The growth rates are measured over three-year periods to smooth out annual noise in employment variation, and because retirement decisions are hardly based on yearly variation. Following the job and worker flow literature, we define the growth rate as $\Delta e_{t,s} / \bar{e}_{t,s}$, where $\Delta e_{t,s}$ is the employment change from year t to year s in a given establishment, and $\bar{e}_{t,s}$ is the average employment level in years t and s . The growth rates defined in this way can take values on the interval $[-2, 2]$. Employment is measured by the number of employees in the establishment at the end of the year. Since employees at risk in year t were by construction all still employed at the end of year $t-1$, the past growth rate is computed over the period from year $t-4$ to year $t-1$. With the current growth rate we refer to the employment change from year $t-1$ to year $t+2$, as the possible transition to disability retirement must take place by the end of year $t+2$.

We believe that labour demand conditions are best described by employment changes in the establishment for those who are employed in large firms with multiple establishments. Therefore, we consider employment variation at establishment level, even though we control the size of the personnel at firm level. In Table 4.2, employees are further divided into groups according to the past and current growth rate of the establishment at which they worked at the end of year $t-1$. Establishments whose growth rate lies on the interval $[-0.10, 0.15)$ are labelled to be 'stable' as opposed to contracting and expanding ones. The average growth during the period under investigation was slightly below 0.05, so that the contracting and expanding establishments are defined in comparison to the growth trend.

First of all, note that our finding that the transition rates to sick leave and disability retirement are highest for the employees of firms with over 1000 employees holds also when we are conditioning on the past or current employment growth category. Somewhat surprisingly, none of the transition rates does seem to be sensitive with respect to the past or current growth rates. Differences in the transition rates between employees at contracting, expanding and stable establishments are generally very small, and do not exhibit consistent patterns. These results should not be taken as conclusive, however, since it is possible that compositional differences in the older workforce across firms of different size, and establishments in different growth categories mask the underlying true relationships.

4.5 Determinants of transition rates

We apply pooled-data logit models to study the determinants of the transition rates. The results from our baseline specification for individual-specific and employer-specific covariates are shown in Table 4.3 and Table 4.4, respectively. We report the odds ratios for dummy variables and the marginal effects for all covariates along with standard errors that are robust to clustering to account for correlation across individuals working in the same establishment. The marginal effects are computed for an average person of the risk group.⁵⁹ We begin our discussion with the impact of individual background characteristics. Then we proceed to the parameters of primary interest, describing the effects of the strictness of medical criteria, experience rating and growth rates.

4.5.1 Individual characteristics

As seen in Table 4.3, the transition rate to sick leave and disability retirement increases strongly with age. For example, the likelihood of entering sick leave is 3.6 percentage points higher at age 62 than at age 50. The difference in the likelihood of becoming disabled grows faster with age and is almost six percentage points higher at 62. These are relatively large increases as the general level of the transition rates is rather low. The average probabilities of entries into sick leave and disability retirement are around 2.2% and 1.9% per year, respectively, as shown in Table 4.3. In terms of the odds ratios, the effect of age on the transition rate from sick leave to disability retirement is largely similar to the effect on the transition rate from work to sick leave.

Women are less likely to move into sick leave and more likely to return to work from sick leave, leading to a lower transition rate to disability retirement. This finding is in line with women's lower incidence of disability in Figure 4.2. The likelihood of being granted a disability pension decreases uniformly with education. The odds of becoming a disability pension recipient is 0.45 for an employee with a Master's degree or higher compared with an otherwise similar employee with a basic education, which corresponds to a 1.2 percentage point lower annual risk of becoming disabled. Education has no effect on the likelihood of returning to work from sick leave. Hence, the lower risk of disability for the educated people is explained by their lower transition rates to sick leave.

⁵⁹ The risk group for transitions from sick leave to disability retirement is different from that for other transitions, and thereby the marginal effects are evaluated at the different values of the covariates. This does not, however, alter our interpretation of covariate effects.

Table 4.3 Odds ratios and marginal effects for individual-specific covariates from baseline logit models

	Probability of a transition to					
	Disability		Sick leave		Db from sick	
	Odds	ME	Odds	ME	Odds	ME
Relaxed medical criteria ^{a)}	1.6623	0.0102	1.6169	0.0116	1.7359	0.0980
	(0.0558)	(0.0009)	(0.0526)	(0.0010)	(0.1310)	(0.0122)
1(sickness benefits > 0 at $t-2$)	2.1834	0.0207	2.3665	0.0272	1.0112	0.0021
	(0.0865)	(0.0015)	(0.0780)	(0.0015)	(0.0769)	(0.0140)
Sickness benefit share ^{b)} at $t-2$		0.0172		0.0126		0.1120
		(0.0042)		(0.0046)		(0.0932)
1(sickness benefits > 0 at $t-3$)	1.9276	0.0166	2.0315	0.0210	1.0359	0.0065
	(0.0654)	(0.0012)	(0.0602)	(0.0012)	(0.0664)	(0.0117)
1(sickness benefits > 0 at $t-4$)	1.7393	0.0134	1.8566	0.0177	0.9806	-0.0036
	(0.0678)	(0.0012)	(0.0615)	(0.0012)	(0.0645)	(0.0122)
Age 51	1.1286	0.0025	1.0419	0.0009	1.0981	0.0171
	(0.0536)	(0.0010)	(0.0384)	(0.0009)	(0.0868)	(0.0142)
Age 52	1.4233	0.0079	1.1866	0.0041	1.2833	0.0447
	(0.0645)	(0.0012)	(0.0429)	(0.0009)	(0.0990)	(0.0133)
Age 53	1.7575	0.0137	1.3616	0.0078	1.4831	0.0695
	(0.0766)	(0.0014)	(0.0478)	(0.0010)	(0.1147)	(0.0128)
Age 54	2.1906	0.0207	1.6736	0.0140	1.3776	0.0569
	(0.0961)	(0.0016)	(0.0604)	(0.0012)	(0.1065)	(0.0131)
Age 55	2.5413	0.0263	1.8173	0.0168	1.5441	0.0758
	(0.1102)	(0.0018)	(0.0652)	(0.0013)	(0.1209)	(0.0127)
Age 56	2.8782	0.0317	1.8498	0.0175	1.5959	0.0810
	(0.1270)	(0.0020)	(0.0709)	(0.0014)	(0.1309)	(0.0132)
Age 57	2.9619	0.0329	1.9254	0.0190	1.6309	0.0844
	(0.1523)	(0.0024)	(0.0827)	(0.0016)	(0.1547)	(0.0150)
Age 58	3.1944	0.0368	1.9852	0.0203	1.7368	0.0941
	(0.1715)	(0.0027)	(0.0914)	(0.0018)	(0.1964)	(0.0173)
Age 59	4.1360	0.0508	2.2305	0.0252	2.1010	0.1223
	(0.2288)	(0.0034)	(0.1180)	(0.0023)	(0.2737)	(0.0185)
Age 60	4.4331	0.0558	2.3901	0.0284	2.4589	0.1436
	(0.2580)	(0.0038)	(0.1306)	(0.0025)	(0.3418)	(0.0182)
Age 61	4.8180	0.0618	2.4240	0.0292	2.6995	0.1553
	(0.2894)	(0.0042)	(0.1403)	(0.0027)	(0.3982)	(0.0184)
Age 62	4.4704	0.0576	2.7432	0.0355	2.4032	0.1399
	(0.2967)	(0.0044)	(0.1679)	(0.0032)	(0.3806)	(0.0210)

Continued on the next page →

	Probability of a transition to					
	Disability		Sick leave		Db from sick	
	Odds	ME	Odds	ME	Odds	ME
Female	0.7456	-0.0056	0.8232	-0.0043	0.6938	-0.0685
	(0.0164)	(0.0004)	(0.0160)	(0.0004)	(0.0343)	(0.0098)
High school	0.9007	-0.0021	0.9157	-0.0021	0.9176	-0.0159
	(0.0173)	(0.0004)	(0.0158)	(0.0004)	(0.0366)	(0.0075)
Lowest tertiary	0.6686	-0.0073	0.6049	-0.0101	1.0299	0.0054
	(0.0226)	(0.0005)	(0.0189)	(0.0005)	(0.0723)	(0.0127)
Bachelor	0.5112	-0.0108	0.4553	-0.0140	0.8765	-0.0245
	(0.0262)	(0.0006)	(0.0224)	(0.0006)	(0.1069)	(0.0231)
Master's or higher	0.4494	-0.0122	0.4061	-0.0153	0.8670	-0.0266
	(0.0335)	(0.0008)	(0.0282)	(0.0008)	(0.1606)	(0.0352)
Foreign language	0.5278	-0.0094	0.7669	-0.0053	0.6269	<i>-0.0909</i>
	(0.0764)	(0.0016)	(0.0844)	(0.0019)	(0.1579)	(0.0512)
Wage position ^{c)}		-0.0087		-0.0061		-0.1186
		(0.0010)		(0.0010)		(0.0186)
Log(earnings)		-0.0005		-0.0055		0.0901
		(0.0008)		(0.0008)		(0.0178)
Spouse working	0.9815	-0.0003	0.9467	-0.0012	1.0859	0.0153
	(0.0181)	(0.0003)	(0.0160)	(0.0004)	(0.0430)	(0.0073)
Spouse retired	1.3619	0.0066	1.2696	0.0059	1.2539	0.0413
	(0.0313)	(0.0006)	(0.0278)	(0.0006)	(0.0688)	(0.0097)
Tenure		0.0001		0.0000		0.0017
		(0.0000)		(0.0000)		(0.0004)

Notes: The models also include controls for year, industry and living region. The reference employee is a 50-year-old single man who has completed only basic education, speaks Finnish or Swedish as his mother tongue, and worked at the end of year $t-1$ in an establishment with 10 to 50 employees. Robust standard errors in (parentheses). Significantly non-unit odds ratios and non-zero marginal effects in bold (95%-confidence level) or in italics (90%-confidence level). a) Dummy for those who may be granted an IER pension by the end of year t . b) The share of sickness benefits of taxable labour income. c) Employee's position in the establishment's wage distribution, ranging from 0 for the lowest wage to 1 for the highest wage.

Employees holding better jobs at the workplace, as measured by their position in the wage distribution, have a lower risk of disability retirement because they are less likely to end up on sick leave and more likely to return to work from sick leave. Family background also matters. Compared with singles, employees whose spouse is still working have a slightly lower probability of sick leave but a higher probability of moving from sick leave to a disability pension. It appears that employees with a retired spouse are the most likely to enter sick leave and the least likely to return to work from sick leave. This may indicate that employees with a retired spouse value their leisure time more than the other groups. Alternatively, their spouses may require special attention at home if they suffer from health problems.

Unfortunately, our data do not contain direct measures of health. To approximate health history we exploit information on the amount of sickness benefits collected in the past years. By construction, the employee at risk in year t did not receive sickness benefits in year $t - 1$. For year $t - 2$ we add a dummy variable indicating whether the employee received sickness benefits, as well as the share of sickness benefits of taxable labour income as a proxy for the fraction of the time spent on sick leave during that year.⁶⁰ For the next two years, we also add dummy variables indicating receipt of sickness benefits. Not surprisingly, receipt of sickness benefits in the past increases the transition rates to sick leave and to disability retirement. Having been on sick leave in year $t - 2$ raises the likelihood of becoming disabled at least by 2.1 percentage points, the overall effect depending on the time spent on sickness benefits. Conditional on being on sick leave, past sickness history has no effect on the likelihood of being granted a disability pension.

4.5.2 Strictness of medical criteria

In 1994 the age threshold for the IER retirement scheme was increased from 55 to 58 for workers born in 1940 or later. In 2000 the entire scheme was effectively abolished from all private-sector employees born in 1944 or later. Given a relatively low emphasis on medical factors when determining eligibility for IER pensions, these reforms can be viewed as increases in the stringency of medical screening for the disability status. It is worth emphasizing that the oldest affected employees were below the pre-reform age thresholds at the time of the reforms. This rules out anticipation behaviour towards IER pensions, providing us with a quasi-experimental setting for evaluating the impact of the medical criteria. More specifically, we exploit the changes in the criteria by including a time-varying dummy variable that equals one for employees born before 1940 who were at least 54 years old at the end of year $t - 1$ and for those born between 1940 and 1943 who were at least 57 years old at the end of year $t - 1$ (*Relaxed medical criteria* in Table 4.3). These groups of employees can potentially qualify for IER pension benefits by the end of year t , so that their disability pension applications are subject to the less strict medical assessment compared with all other employees.

Being eligible to apply for an IER pension clearly raises all three transition rates. The odds of entering sick leave is 1.6, implying a 1.2 percentage point higher transition rate to sick leave for an employee who can apply for an IER pension benefit than for an otherwise similar non-eligible employee. Conditional on being

⁶⁰ This is not an accurate measure because the waiting time until the receipt of sickness benefits from KELA can vary between employees, and because the amount of a sickness benefit is determined as a decreasing fraction of the past earnings.

on sickness benefits, the IER pension scheme increases the likelihood of being granted a disability pension by 9.8 percentage points. The overall probability of moving to the disability pension track is one percentage point higher for employees with an option to apply for an IER pension. Not surprisingly, these estimates are somewhat sensitive with respect to the treatment of workers who moved directly from work to disability retirement without receiving sickness benefits first. If these workers are removed from the pool of sickness benefit recipients, the impact of IER scheme eligibility becomes weaker: its effect on the odds of entering sick leave drops to 1.3 (and the associated marginal effect to 0.0046) and that on the odds of moving from a sick leave to disability retirement reduces to 1.2 (and the marginal effect to 0.0327). Nevertheless, all the effects remain statistically significant at the conventional confidence levels, implying that our qualitative results are robust. Overall, our findings are in accordance with Börsch-Supan's (2007) conclusion that the strictness by which vocational considerations (at the expense of medical criteria) are applied when determining eligibility for disability pension benefits is strongly related to disability pension incidence.

4.5.3 Experience rating

Next we turn to the effects of the covariates that are closely related to the experience-rated contributions: firm size, disability cost and equity ratio. The first determines the degree of experience-rating but could have an effect on its own. The second measures the expected lump-sum payment the employer has to pay in the case the employee is granted a disability pension in year t . This marginal cost of the disability event is a function of the worker's age, disability pension benefit (determined by earnings history) and the degree of experience-rating applied to the employer (determined by firm size in a given year).⁶¹ As we control for age, earnings, firm size and year fixed effects in the model, the effect of disability cost is identified by the two reforms in the experience rating scheme that were described in section 3.2. The third variable, the equity ratio, measures the firm's ability to incur disability pension costs.

The model includes three dummy variables for the size of the employing firm at the end of year $t - 1$. The smallest firms with 10 to 50 employees, which were not subject to the experience-rated contributions in any year, serve as the reference category. This size categorization is relatively coarse, but our results are not sensitive with respect to different specifications of the firm size effects. As seen in Table 4.4, the likelihood of a transition from work to sick leave increases with

⁶¹ The disability costs are estimated using the formula in equation (1). Since the true level of the disability pension benefit is not known, we assume the pension benefit would be a fixed per cent (55%) of the annual earnings.

firm size, being significantly higher for the largest firms. Namely, an employee of a firm with over 1000 employees has a 0.9 percentage point higher risk of entering sick leave per year than an otherwise similar person in a firm with 50 or fewer employees.

Table 4.4 Odds ratios and marginal effects for employer-specific covariates from baseline logit models

	Probability of a transition to					
	Disability		Sick leave		Db from sick	
	Odds	ME	Odds	ME	Odds	ME
Log(disability cost)		-0.0007		-0.0006		-0.0106
		(0.0001)		(0.0001)		(0.0022)
Equity ratio ^{a)}		0.0037		-0.0009		0.0948
		(0.0012)		(0.0011)		(0.0196)
Past employment growth from $t - 4$ to $t - 1$		0.0003		0.0000		-0.0018
		(0.0005)		(0.0005)		(0.0102)
Current employment growth from $t - 1$ to $t + 2$		-0.0007		-0.0007		0.0115
		(0.0005)		(0.0005)		(0.0101)
Past excess worker turnover from $t - 4$ to $t - 1$		0.0000		0.0000		-0.0224
		(0.0012)		(0.0014)		(0.0237)
Current excess worker turnover from $t - 1$ to $t + 2$		0.0046		0.0040		0.0812
		(0.0012)		(0.0012)		(0.0244)
Log(turnover)		0.0009		0.0005		0.0132
		(0.0003)		(0.0003)		(0.0049)
Establishment closure at t	0.4327	-0.0113	0.5405	-0.0105	0.4989	-0.1380
	(0.0571)	(0.0012)	(0.0676)	(0.0016)	(0.1410)	(0.0589)
Firm size 51–300	1.1753	0.0028	1.2396	0.0042	<i>1.2889</i>	0.0470
	(0.0644)	(0.0010)	(0.0613)	(0.0011)	(0.1504)	(0.0209)
Firm size 301–1000	1.1408	0.0023	1.2847	0.0049	1.1288	0.0228
	(0.0950)	(0.0015)	(0.0926)	(0.0016)	(0.1836)	(0.0302)
Firm size > 1000	1.3371	0.0054	1.5206	0.0089	1.0971	0.0175
	(0.1249)	(0.0020)	(0.1229)	(0.0021)	(0.1988)	(0.0339)
Mean age of employees		0.0000		-0.0001		0.0011
		(0.0001)		(0.0001)		(0.0012)
N observations	887,321		887,321		19,858	
N establishments	28,961		28,961		7,438	
Pseudo R^2	0.085		0.079		0.219	

Notes: The models also include controls for year, industry and living region. The reference employee is a 50-year-old single man who has completed only basic education, speaks Finnish or Swedish as his mother tongue, and worked at the end of year $t - 1$ in an establishment with 10 to 50 employees. Robust standard errors in (parentheses). Significantly non-unit odds ratios and non-zero marginal effects in bold (95%-confidence level) or in italics (90%-confidence level). a) Equity as share of assets, bottom coded below the 5th percentile and top coded above 95th percentile.

In the large firms many people are doing similar work, making it easier to share job tasks of a sick person between the remaining employees. This may induce the employees of large firms to apply for sickness benefits more frequently, which could explain our finding. On the other hand, the likelihood of disability retirement conditional on being a recipient of sickness benefits is almost independent of the firm size. As a consequence, there is a positive relationship between the firm size and the overall risk of ending up with a disability pension. Note that these estimates should describe the true firm size effects, as the disability cost variable accounts for the effect of the degree of experience-rating.⁶²

The expected cost of disability retirement has a negative effect on all the transition rates (see Table 4.4). These effects are accurately estimated, but the magnitude of the marginal effects is very small. For a recipient of sickness benefits a ten percentage point increase in the disability cost decreases the likelihood of being awarded a disability pension by 0.1 percentage points. The other two marginal effects are even smaller. However, one should note that the disability cost variable exhibits a large degree of variation, ranging from zero to several times the annual earnings. To get a better picture of the effects of experience-rating, we computed the transition probabilities at different ages and different values of the disability cost variable, holding all the other covariates fixed at their sample means. In the absence of the disability cost due to experience rating, the likelihood of being granted a disability pension at age 55 would be 0.024.

When the disability cost are introduced and set to the median value of disability costs in the sample, this probability declines to 0.017, i.e. a decrease of about 30%. Furthermore, by introducing the maximum degree of experience-rating, it is possible to obtain some 50% decline in the disability risk at age 50 compared with the case of no experience rating. A somewhat larger part of these experience rating effects can be attributed to the decline in the transition rate to sick leave, but the increase in the likelihood of returning to employment from sick leave plays a notable role as well. In other words, both the *ex ante* and *ex post* effects of experience-rating are not only statistically significant, but also economically important.⁶³

⁶² Should we exclude the disability cost from the model, there would be a large effect of firm size on the transition rate from sick leave to disability retirement. For the employees of the two largest employer groups, the odds of moving from sick leave to disability retirement would be about 0.66, corresponding to a decrease of 9 percentage points in the disability probability for the average recipient of sickness benefits.

⁶³ It is not clear at which point the employers became aware of the new experience-rating scheme before the law changes. The reforms may also have induced some sort of anticipatory behaviour just before the

Since experience rating aims to affect employer behaviour through financial incentives, the effectiveness of such incentives should depend on the firm's financial position. When the experience-rated firm is short of liquid assets, it might try harder to deter exits to disability retirement to avoid the costs that in a dire financial situation might bankrupt the firm. This is a relevant concern especially in the Finnish system where the employer has to pay its share of the present value of disability costs as a lump-sum payment at the time when a disability pension is granted to its employee. To address this question we use the equity ratio as a proxy for the firm's financial position.⁶⁴ In Table 4.4, the equity ratio has a positive effect on the likelihood of being granted a disability pension. It also has a strong effect on the likelihood that a sick leave will be followed by a disability pension. This implies that a recipient of sickness benefits in a firm in a weak financial position returns to work with a relatively high probability.

Given that the experience-rated contributions depend on firm size and that larger firms may have better possibilities to organize retraining and arrange alternative job tasks for their employees with reduced working capacity, we should expect the effect of the equity ratio to vary across firms of different size. We therefore extend our baseline specification by adding interaction terms of firm size categories and equity ratio. The results of this exercise are shown in Table 4.5.

Table 4.5 *Marginal effects for the equity ratio by firm size from extended logit models*

	Probability of a transition to		
	Disability	Sick leave	Db from sick
Equity ratio ^{a)} for group			
Firm size 10–50	-0.0017 (0.0020)	-0.0030 (0.0021)	0.0470 (0.0362)
Firm size 51–300	<i>0.0030</i> (0.0016)	0.0010 (0.0016)	0.0658 (0.0311)
Firm size 301–1000	0.0077 (0.0030)	0.0034 (0.0030)	0.0278 (0.0520)
Firm size > 1000	0.0068 (0.0025)	-0.0033 (0.0024)	0.1921 (0.0415)

Notes: Other covariates as in the baseline specification. Robust standard errors in (parentheses). Significantly non-zero marginal effects in bold (95%-confidence level) or in italics (90%-confidence level). a) Equity as share of assets, bottom coded below the 5th percentile and top coded above 95th percentile.

new rules came into effect. However, if we drop one or two years preceding the reforms from the analysis, our results do not change notably.

⁶⁴ There are a few extreme values for the equity ratio. To deal with such outliers we bottom and top coded the equity ratio at the 5th and 95th percentiles, respectively. In other words, we use the threshold values for observations below the 5th percentile or above the 95th percentile.

There is no evidence of statistically significant effects in the smallest firms that are not subject to experience rating, which is consistent with the claim that the effect of the equity ratio is attributable to the experience-rating system. In firms that employ more than 300 employees, the likelihood of being granted a disability pension increases with the equity ratio. In the case of the largest firms, the likelihood of being granted a disability pension for a recipient of sickness benefits increases with the equity ratio: an increase of ten percentage points in the equity ratio is related to a 1.9 percentage point increase in the disability pension incidence. Put differently, large firms that can afford the cost of disability pension expenditures seem to put less effort into occupational rehabilitation compared with large firms in a weaker economic position.

In sum, our findings give strong support for the hypothesis that experience rating affects employer behaviour. The higher expected cost of the disability event for the employer lowers transition rates from work to sick leave and from sick leave to disability retirement. This suggests that the firms subject to experience rating apply preventive measures to minimize entries to sick leave, as well as put more effort in occupational rehabilitation to get their employees on sickness benefits back to work. Our results for the effect of the equity ratio give further support for the importance of experience-rating effects. Namely, the financial position of a firm has an effect on the employees of the larger firms that are liable for a significant fraction of disability pension expenditure of their former employees, whereas we find no relationship between the equity ratio and transition rates in the firms that are not subject to experience rating.

4.5.4 Employment growth and excess turnover

In addition to the growth rates, our models include a control variable for excess worker turnover, which measures the degree of restructuring at the workplace for a given net change in employment.⁶⁵ In general, the growth and excess turnover rates are affected by the outcome of interest. In other words, the worker's transition out of work. To eliminate the resulting endogeneity problem, we have adjusted these covariates for each worker by removing the effect of the worker's own mobility in and out of the establishment.

⁶⁵ Excess worker turnover in year t is defined as $h_t + s_t - |\Delta e_t|$, where h_t and s_t denote the number of hires and separations during year t , respectively, and Δe_t is the employment change from the end of year $t-1$ to the end of year t in a given establishment. This quantity is the worker flow in excess of what is needed to explain the net change in the size of establishment's workforce. Dividing it by the average employment level at the end of years $t-1$ and t , say \bar{e}_t , gives the excess turnover rate: $(h_t + s_t - |\Delta e_t|) / \bar{e}_t$, which takes a value on the interval $[0, 2]$. To smooth annual variation we take the average of the excess turnover rates between years $t-4$ and $t-1$ (for the past period) and between years $t-1$ and $t+2$ (for the current period).

Excess worker turnover can be either voluntary or involuntary from the employees' standpoint. It may result from the restructuring measures through which the employer adjusts the structure of the workforce. Or it may be driven by a high level of voluntary quits, perhaps induced by poor working conditions, management or wage rates, which are compensated by new hires. As seen in Table 4.6, excess worker turnover over the past three years has no effect on the transition rates. However, employees in the establishments with high current levels of excess worker turnover are more likely to enter sick leave and less likely to return to work from sick leave. Consequently, exits via disability retirement are more common in workplaces with a high rate of excess worker turnover. It should be noted that the effect of excess worker turnover is conditional on a given change in the employment level, as we control for the employment growth rates. High turnover can result in extra training work for the tenured employees and cause other problems at the workplace, and thereby lead to an increase in stress factors. When high turnover reflects an ongoing restructuring process, our estimates suggest the possibility that the employer encourages some older employees to apply for a disability pension. For an employer that is adjusting the structure of the workforce but is not downsizing, such a policy can be an effective alternative for dismissals that would be difficult to justify. If so, we should expect to find a positive effect for the employment growth rate as well, but none of the effects of the employment changes during the three-year periods differs statistically significantly from zero in Table 4.4.

The underlying assumption of symmetric effects for the expansion and contraction of the workforce is quite restrictive. In Table 4.6 we therefore report results from model specifications that do not impose such a restriction but allow for different coefficients for positive and negative growth rates. It seems that the risk of being granted a disability pension is lower for employees holding jobs in establishments that are currently either downsizing or expanding. The effects of the current growth rates can be attributed to the increased risk of sick leave, whereas the transition rate from sick leave to disability retirement is not affected by the current growth rates.

Table 4.6 Marginal effects for employment growth rates from extended logit models

	Probability of a transition to		
	Disability	Sick leave	Db from sick
Past growth $t - 4$ to $t - 1$			
when ≥ 0	0.0018 (0.0006)	0.0008 (0.0006)	0.0184 (0.0132)
when ≤ 0	-0.0039 (0.0012)	-0.0021 (0.0013)	-0.0628 (0.0250)
Current growth $t - 1$ to $t + 2$			
when ≥ 0	-0.0187 (0.0018)	-0.0190 (0.0019)	0.0094 (0.0321)
when ≤ 0	0.0030 (0.0007)	0.0029 (0.0007)	0.0152 (0.0123)

Notes: Other covariates as in the baseline specification. Robust standard errors in (parentheses). Significantly (95%-confidence level) non-zero marginal effects in bold. 90%-confidence level indicated with italics.

The marginal effect of the current decline in employment on the probability of disability pension receipt is 0.003. For example, a 50% decrease in the workforce over the next three years (i.e. the growth rate of -0.67) is estimated to reduce the disability pension entry rate by 0.2 percentage points ($= -0.67 \times 0.003 \times 100$) compared with the case of no change in the size of the workforce. One possible explanation is that employees with health problems are less willing to apply for sick leave at the times when they are worried about their jobs. During slack demand employers may also use dismissals to get rid of employees with reduced working capacity before they apply for a sickness benefit or disability pension. In any case, the effect of downsizing is very small. It also implies that the employers do not encourage early retirement through the disability schemes as a soft way of downsizing.

The marginal effect of the current employment expansion on the likelihood of disability pension receipt is quite large, being -0.0187 . Thus, being employed in an establishment whose workforce increases by one-half by the end of year $t + 2$ (i.e. the growth rate of 0.4) decreases the probability of being granted a disability pension by 0.75 percentage points ($= 0.4 \times -0.0187 \times 100$) compared with the case of working in a stable establishment. When a firm is expanding its business rapidly, it may experience difficulties in hiring the sufficient amount of skilled labour. In such a case, the employer may put some extra effort to keep its old employees at work, which may explain our finding. Furthermore, for the employee an expansion period may indicate better economic opportunities, in terms of promotion possibilities or extra pay, which increase the value of staying employed despite some health problems.

Black *et al.* (2002) found that negative (positive) demand shocks increase (decrease) the entry rate to disability benefit schemes in aggregate US data. We have just shown that the reverse relationship holds for the negative shocks at the establishment level in the Finnish labour market. Of course, one should bear in mind that we are considering only transitions out of work. During economic downturns, transitions from nonparticipation and unemployment to disability schemes are likely to increase, and such transitions may dominate the US data.

Compared with the impact of the current growth rates, the past growth rates have the opposite effects on the likelihood of being granted a disability pension. Namely, both the reduction and expansion in the workforce during the past three years increase the transition rate to disability retirement. Hence, our findings are in accordance with the results of Vahtera *et al.* (2005) and Rege *et al.* (2009), who found that a large reduction of the workforce in the past leads to a notable increase in the entry rate to disability retirement. This is quite remarkable given the differences in the research design. Recall that the risk set of Rege *et al.* included also those who lost their job as a result of plant downsizing, whereas Vahtera *et al.* considered only employed workers but their data came from the municipal sector and cover an exceptional period of deep recession.

4.6 Concluding remarks

In this study we analysed how labour demand and institutional factors affect transitions to sick leave and disability retirement. Using matched employer-employee data for the Finnish private sector, we were able to measure the employment growth rates and excess worker turnover at the establishment level. To study the role of the institutional setting, we exploited the law changes that affected the medical requirements for disability pension eligibility and the partially experience-rated employer contributions. Our main findings can be summarized as follows:

- For older employees a transition to sick leave is often a one-way street out of employment, leading eventually to disability retirement. Roughly half of 50–55 year-olds and over two-thirds of older workers on sickness benefits end up in disability retirement within the next three years. This highlights the importance of preventive measures aimed at minimizing the flow into sick leave.
- Those employees who can apply for a disability pension under more lenient medical requirements are much more likely to enter sick leave and to retire via disability pension benefits. Therefore, the abolition of the individual early retirement scheme in 2000 significantly reduced the flow into disability retirement in the affected groups.

- There is ample evidence that experience rating lowers the flow into sick leave (i.e. the *ex ante* effect), and reduces transitions from sick leave to disability retirement (i.e. the *ex post* effect). Moreover, large firms that can easily bear their share of early retirement costs due to their strong financial position let their employees on sickness benefits exit more easily via disability pension schemes than firms in a weaker position do. Financial situation is not an issue for smaller firms that are not subject to the experience rating.
- The transition rates to sick leave and disability retirement are relatively large in establishments experiencing a high degree of excess worker turnover. When an establishment is growing, transitions to sick leave and disability retirement become less frequent. There is no evidence of employers exploiting the disability pension scheme as a way of adjusting their workforce when downsizing.

These findings imply two policy recommendations to reduce the disability benefit enrolment rate of older workers. First, the stringency of medical criteria and medical screening for disability benefit eligibility should be tough enough. When non-medical factors are weighted at the expense of medical criteria, disability benefits may distort labour supply decisions, thereby also inducing workers who are not truly disabled to retire via disability programmes. This appears to be mainly a labour supply issue, as we did not find evidence of employers encouraging disability retirement when downsizing. Secondly, the experience-rating of disability benefit costs seems to be an effective policy instrument. It seems to induce employers to take preventive actions to reduce the inflow into sick leave, and to put more effort into get their employees on sickness benefits back into work. This finding should be of considerable interest, not only for Finland, but also for other countries that do not have an experience-rating system for disability benefits (yet). Obviously, there are still a number of open questions, regarding, for example, the optimal design of experience-rating and possible spillover effects on hiring and transitions out of work to other destinations than disability retirement. These questions need to be addressed in order to get a more complete picture of the consequences of the experience-rating of disability benefits.

References

- Autor, H. D. – Duggan, M. G. (2003): The Rise in the Disability Rolls and the Decline in Unemployment. *Quarterly Journal of Economics* 118(1):157–205.
- Autor, H. D. – Duggan, M. G. (2006): The Growth in the Social Security Disability Rolls: A Fiscal Crisis Unfolding. *Journal of Economic Perspectives* 20(3):71–96.

- Black, D. – Daniel, K. – Sanders, S. (2002): The Impact of Economic Conditions on Participation in Disability Programmes: Evidence from the Coal Boom and Bust. *The American Economic Review* 92(1):27–50.
- Börsch-Supan, A. (2007): Work Disability, Health, and Incentive Effects. MEA Discussion Paper 135–2007, Mannheim.
- Campolieti, M. (2004): Disability Insurance Benefits and Labor Supply: Some Additional Evidence. *Journal of Labor Economics* 22(4):863–888.
- Gielen, A. C. – Van Ours, J. C. (2006): Age-Specific Cyclical Effects in Job Reallocation and Labor Mobility. *Labour Economics* 13(4):493–504.
- Gruber J. (2000): Disability Insurance Benefits and Labor Supply. *The Journal of Political Economy* 108(6):1162–1183.
- Hassink, W. H. J. – Van Ours, J. C. – Ridder, G. (1997): Dismissal Through Disability. *De Economist* 145(1):29–46.
- Hutchens, R. (1999): Social Security Benefits and Employer Behaviour: Evaluating Social Security Early Retirement Benefits as a Form of Unemployment Insurance. *International Economic Review* 40(3):659–678.
- Koning, P. (2009): Experience Rating and the Inflow into Disability Insurance. *De Economist* 157(3):315–335.
- Kyyrä, T. – Wilke, R. (2007): Reduction in the Long-Term Unemployment of the Elderly: A Success Story from Finland. *Journal of the European Economic Association* 5(1):154–182.
- McVicar, D. (2008): Why Have UK Disability Benefit Rolls Grown So Much? *Journal of Economic Surveys* 22(1):114–139.
- OECD (2008): *Sickness, Disability and Work: Breaking the Barriers*. Vol. 3: Denmark, Finland, Ireland and The Netherlands. Paris.
- Rege, M. – Telle, K. – Votruba, M. (2009): The Effect of Plant Downsizing on Disability Pension Utilization. *Journal of the European Economic Association* 7(4):754–785.
- Saurama, L. (2004): Experience of Early Exit: A Comparative Study of the Reasons for and Consequences of Early Retirement in Finland and Denmark in 1999–2000. Finnish Centre for Pensions Studies, 2004:2.
- Vahtera, J. – Kivimäki, M. – Forma, P. – Vikström, J. – Halmeenmäki, T. – Linna, A. – Pentti, J. (2005): Organisational Downsizing as a Predictor of Disability Pension: The 10-Town Prospective Cohort Study. *Journal of Epidemiology & Community Health* 59(3): 238–242.

5. Employment and wage effects of a payroll tax cut – evidence from a regional experiment⁶⁶

Abstract

In this paper, we evaluate the effects of a regional experiment that reduced payroll taxes by 3–6 percentage points for three years in northern Finland. We match each firm in the target region with a similar firm in a comparison region and estimate the effect of the payroll tax reduction by comparing employment and wage changes within the matched pairs before and after the start of the experiment. According to our results, the reduction in payroll taxes led to an increase in wages in the target region. The point estimates indicate that the increase in wages offset roughly half of the impact of the payroll tax cut on labour costs. The remaining labour cost reduction had no significant effects on employment.

Keywords: Payroll tax, Labor demand, Tax incidence, Propensity score matching

JEL Classification J18, J23, J38, J58, J65, J68

5.1 Introduction

A reduction in payroll taxes lowers wage costs and hence boosts the demand for labour. Its effect on employment depends on the incidence of the payroll taxes. If the tax cut leads to higher wages that entirely offset the reduction in taxes, the tax cut has no effect on employment.

Past evidence on the incidence and the employment effects of payroll tax changes is mixed. Studies that rely on cross-country or time-series variation in national payroll taxes produce widely varying estimates of tax incidence. An important problem in such approaches is the omitted variables bias. In cross-country studies, it is difficult to control for all the differences in wage-setting institutions. These unobserved, across-country differences may be correlated with differences in the level of taxation and employment. In the time-series studies, there may be simultaneous changes in other variables that affect wages and employment. For

⁶⁶ Ossi Korkeamäki and Roope Uusitalo.

This research project was funded by the Finnish Ministry of Social Affairs and Health. The authors would like to thank Merja Seppä-Heikka and Seppo Ritari from the National Board of Taxes for their help with the data and Eva Mörk, Erik Mellander and Hans Lind and the seminar participants at Uppsala, Amsterdam, San Francisco and Helsinki for their helpful comments.

example, Hamermesh (1993) summarizes this literature by noting, “The estimates of tax shifting vary across the entire admissible range and even outside it” and concludes, “It is impossible to draw any firm conclusions about the incidence of payroll tax from these studies”.

A more promising approach is to examine the effects of changes in taxes or other mandatory employer contributions when these changes differ across otherwise similar firms. Following this approach, Gruber (1994) evaluates the effects of mandated maternity benefits in the US, and Gruber (1997) the effects of changes in mandatory pension contributions in Chile. Anderson and Meyer (1997) and Murphy (2007) examine the incidence of unemployment insurance taxes in the US. In all these cases, the changes in the payroll tax rates vary between firms because of the different composition of their labour force or because the tax rates depend on firm characteristics. Another approach that is more directly related to our study examines the effects of regional policies that create different changes in the payroll tax rates across firms that are located in different regions but that are otherwise comparable. Prime examples include Bohm and Lind (1993), who evaluate the employment effects of regional wage subsidies in northern Sweden, Johansen and Klette (1998), who examine the effects of regional differences in payroll taxes in Norway, and Bennmarker, Mellander and Öckert (2009), who evaluate the effects of a recent regional wage subsidy scheme in Sweden. These studies typically find that changes in payroll taxes are mostly shifted to wages with little effect on labour costs or employment.

In this paper, we evaluate the employment and wage effects of a regional experiment in northern Finland. This experiment abolished employer contributions to the national pension scheme and the national health insurance for firms located in the targeted high unemployment regions. Prior to 2003, these employer contributions varied between 2.95 and 6 per cent of the wage bill, depending on the capital intensity and size of the firm. From January 1 2003, all private employers in the 20 target municipalities located in northern Finland and on the islands along the western coast were exempt from these social security contributions for three years. In this paper we focus on the effects in northern Finland, where over 90 per cent of the eligible firms are located.

A regionally targeted programme has several benefits compared to an across-the-board cut in taxes. Perhaps the main benefit for policymakers is that the effects of a regional programme are substantially easier to evaluate. The employment change in the target region can be compared to similar regions that are not affected by the tax cut. If the target and comparison regions are truly similar, estimates of the employment effects based on differences in the employment and wage changes between the treatment and comparison regions provide much more reliable estimates of the effects of the payroll tax cut than time-series or cross-section variation in payroll taxes could ever do.

We use firm-level data to evaluate the effects of the payroll tax cut on employment and individual data to evaluate the effect on wages. Our main results are based on a comparison of the employment changes in target-region firms and the employment changes in firms located in a control region that is as similar as possible in terms of unemployment rate, industry structure and the composition of the labour force. Finally we compare target-region firms in northern Finland to firms located in other high-unemployment areas in northern and eastern Finland.

Comparison of the employment changes across regions still creates problems if the regions are not quite similar in all relevant characteristics. For example, an industry-specific boom might have different effects in different regions depending on the industry structure of the region. To make the treatment and comparison regions more comparable, we adopt a matching procedure to identify comparable firms (or rather plants) in the treatment and control regions. We then evaluate the effects of the payroll tax cut by comparing firms located in different regions but otherwise similar in all their observed pre-treatment characteristics.

5.2 The experiment

Payroll taxes in Finland consist of employer contributions to the employees' pension scheme, the unemployment insurance, the national pension insurance, the national health insurance, and the employment accident insurance. The tax rates of various components vary across sectors and by firm size⁶⁷. According to Statistical Yearbook of the Social Insurance Institution, the average payroll tax rate was 23.86% in 2002.

In March 2002, the Finnish government agreed to a temporary removal of employer contributions to the national pension insurance and the national health insurance for firms that operated in the twenty target municipalities. Removal of these contributions lowered the payroll taxes for the eligible firms by 4.1 percentage points, on average. The programme was designed as an experiment with a stated aim to evaluate the effect of a cut in the payroll taxes on employment in the target region. The payroll tax exemption lasted for three years from January 1 2003 to December 31 2005. In December 2005, the government extended the duration of the experiment to the end of 2009.

⁶⁷ In 2002, the private sector employers contributed 1.69% of the wage bill to the national health insurance, and 1.00% to the employment accident insurance. For calculating the national pension insurance contributions the firms are divided into three categories based on their size and capital intensity. The contribution rates in these categories were 1.35, 3.55 and 4.45. The Unemployment Insurance contributions are progressive, the contribution rate being 0.7% of wage bill for wages up to 840,940 euros and 2.7% of the wages exceeding this threshold. The Employees' Pension Scheme has a relatively complicated fee structure. In the large firms, pension contributions vary with the age of the employee and are partially experience-rated and depend on the number of previous employees receiving early retirement benefits. Small firms pay a flat rate of 17.32%.

As the payroll tax exemption may have anticipatory effects, it is useful to note that the tax exemption was first suggested by a working group that presented its report in December 2001. The law was a part of the government budget proposal for the year 2003 that was agreed upon within the government in March 2002. The government gave the proposal to the parliament in September 2002. The parliament accepted the budget proposal and the president signed the law on the payroll tax exemption in December 2002. The payroll tax exemption was also widely discussed in press during the spring 2002. It is, therefore, possible that firms who anticipated the tax exemption could have altered their employment already before the start of the programme in January 2003. However, it is unlikely that any employment effects could have occurred before March 2002 since the nature of the programme was very much an open question until then.

All private employers and state-owned enterprises that had a “permanent place of business” in the twenty target municipalities were eligible for the tax exemption. The maximum annual reduction was 30,000 euros per firm. To comply with the EU-legislation regulating state-aid that may distort competition within the Union, agriculture, fishing, and transport industries were excluded from the experiment. An important restriction is also that local governments were not eligible for the exemption.

Prior to the beginning of the experiment the government estimated that 3500 firms would be eligible for the exemption, and that the budgetary cost of the experiment would be eight million euros. To cover the costs without cutting benefits financed by payroll taxes the experiment was financed by temporarily raising the national health insurance contributions for the employers outside the target region by 0.014 percentage points.

All the target municipalities were located in high unemployment areas. However, the geographical borders of the target area were somewhat arbitrary. There were other regions outside the target area with comparable, and even higher, unemployment rates. The target municipalities were selected through a political process and there is no obvious reason why just these municipalities were selected. In fact, the original task of the working group that proposed the tax exemption was limited to measures that would be targeted only to the three northernmost municipalities. In their final report, the working group proposed two alternatives: one involving only these three municipalities and another involving nine other municipalities in the northern Finland. After the working group rendered its final report, but before the government gave its proposal to the parliament, two more municipalities in Lapland and six municipalities on islands along the western coast were added to the tax exemption region. Eventually the target area covered the entire province of Lapland except its capital region around Rovaniemi and an industrial region around Kemi-Tornio. On the other hand, the working group would have granted a tax exemption also to the local

government employers. The final proposal was a compromise that excluded all public sector employers with the exception of state-owned enterprises⁶⁸.

Applying for the tax exemption was made easy for the participating employers. The employers were only required to file a starting declaration to the local tax office. The employers could then simply deduct the tax-exempt amount from their monthly employer contributions. An additional requirement was that the employers also had to report tax exemptions in detail in their annual report to the tax administration. The ease of participation was reflected in high take-up rates. According to our calculations, all eligible employers with at least 50 employees, 90 per cent of the eligible employers with at least five employees and 75 per cent of the firms with 2–4 employees had filed a starting declaration by December 2003.

Table 5.1 Participating firms according to size

Firm size (number of full-time employees)	N firms	N Employees	Payroll tax deduction
0	456		31,955
1	424	424	84,075
2–4	659	1836	382,585
5–9	369	2451	686,321
10–19	237	3202	931,020
20–49	139	4153	1,157,750
50–99	37	1544	600,498
101–250	10	1578	289,497
> 250	3	911	63,555
Total	2334	17,099	4,227,256

Source: Authors calculations from data provided by the National Board of Taxes.

Most firms that applied for the tax exemption were very small. The median firm had only four employees. Only ten per cent of the firms had more than twenty, and 2.5 per cent more than fifty employees. In terms of employment and payroll tax bill, these “large” firms naturally represent a much higher share. The largest industries were business services, retail trade, hotels and restaurants, and construction. In total, the experiment involved 2334 firms with 17,099 employees during the first year. According to our calculations, the reduction of payroll tax revenue due to the experiment was 4.2 million euros in the first year.

⁶⁸ State-owned enterprises are government agencies that operate in the market and compete with the private firms. The largest such agencies in the target area are Destia that builds and maintains roads and Metsähallitus that mainly maintains state-owned forests and national parks.

5.3 Tax incidence and the Finnish wage bargaining system

According to the textbook model, the incidence of payroll tax cut depends on the relative elasticities of labour demand and labour supply. A typical empirical finding is that labour supply is relatively inelastic and that the workers therefore bear the cost of tax increases. Most often these incidence results are presented within the context of competitive labour markets. However, the Finnish wage determination system differs substantially from the competitive market model. Below we describe the main features of the system focusing on its potential effects on tax incidence.

Wage bargaining in Finland involves a high degree of co-ordination between the different unions and the employer organizations. A framework agreement is typically negotiated at a national level between the union and employer federations on a one- or two-year basis. After central agreement has been reached, the individual unions and the respective employer organizations bargain over wages separately in each industry. These contracts determine a general wage increase applied to all wages in the sector and a wage schedule determining a minimum pay in each task. The industry-specific collective labour agreements are also binding for the non-union members in the industries where the union contract is “representative”. Since union density is roughly 70 percent, most industries have a representative contract. There are no statutory minimum wages in Finland.

Even though union bargaining occurs at the national level, there is room for regional variation in wages, as well as, wage variation across firms and across workers within firms. Local bargaining between individual workers or their local union organization and the firms may lead to outcomes that deviate from the national contracts. The employer can naturally pay more and if both local parties agree, and as long as the minimum provisions are not violated, even less than what is agreed in the national contract. Wage drift, defined as wage increase exceeding what is agreed in the union contracts, has historically accounted for approximately 40 per cent of the wage growth. This fraction has declined over time but was still 35 per cent between 1992 and 2000 (Uusitalo 2005). More recently firm-specific arrangements such as profit sharing have become more important, leading to an increase in across firm variance in wages (Uusitalo and Vartiainen 2007).

The implications of national wage bargaining for the tax incidence are not entirely clear. On the one hand, one might claim that wage changes are determined at the national level and that a regional payroll tax subsidy scheme has little or no impact on wages. On the other hand, the importance of local bargaining and the fact that the employee and the employer can freely agree on wage increases exceeding what is agreed in the union contract, may lead to a situation where the payroll tax cut leads to a wage increase. Even in this case, tax

shifting may be different than in a national scheme if the workers are mobile and due to factor mobility the elasticity of local labour supply larger than elasticity of labour supply in the whole country (Murphy 2007).

5.4 Empirical strategy

Our estimates are based on differences in employment and wage changes between the firms eligible for tax exemption and a control group. We created the control group by a two-stage procedure. We first selected the “counties” (NUTS4-level sub-regional units) that were most comparable to the target region in terms of unemployment rates, industrial structure and workforce characteristics. We based the selection on the regional statistics published in “Seutukunta- ja maakuntakatsaus 2002” by Statistics Finland. The target region had a high unemployment rate and little manufacturing or other industrial activity. The share employed in agriculture was much higher and the average level of education much lower than in the rest of the country. To create a comparable control region we excluded two regions from the other non-target regions in Lapland because they were administrative centres with above average education level (Rovaniemi) or major manufacturing regions (Kemi-Tornio). Instead, we included high unemployment areas from Eastern Finland just south of the target region. Also in choosing the comparison area, we excluded regions with major cities so that the comparison area would resemble the target area also in its industry composition (see map in Figure 5.1). Our judgment is that the choice of comparison areas was rather successful. As shown in Table 5.2, the target and control regions have similar unemployment and employment rates, reasonably similar industry distribution and a similar population structure. In all these dimensions, the target and control regions deviate substantially from the national average.

Figure 5.1 Target and comparison regions

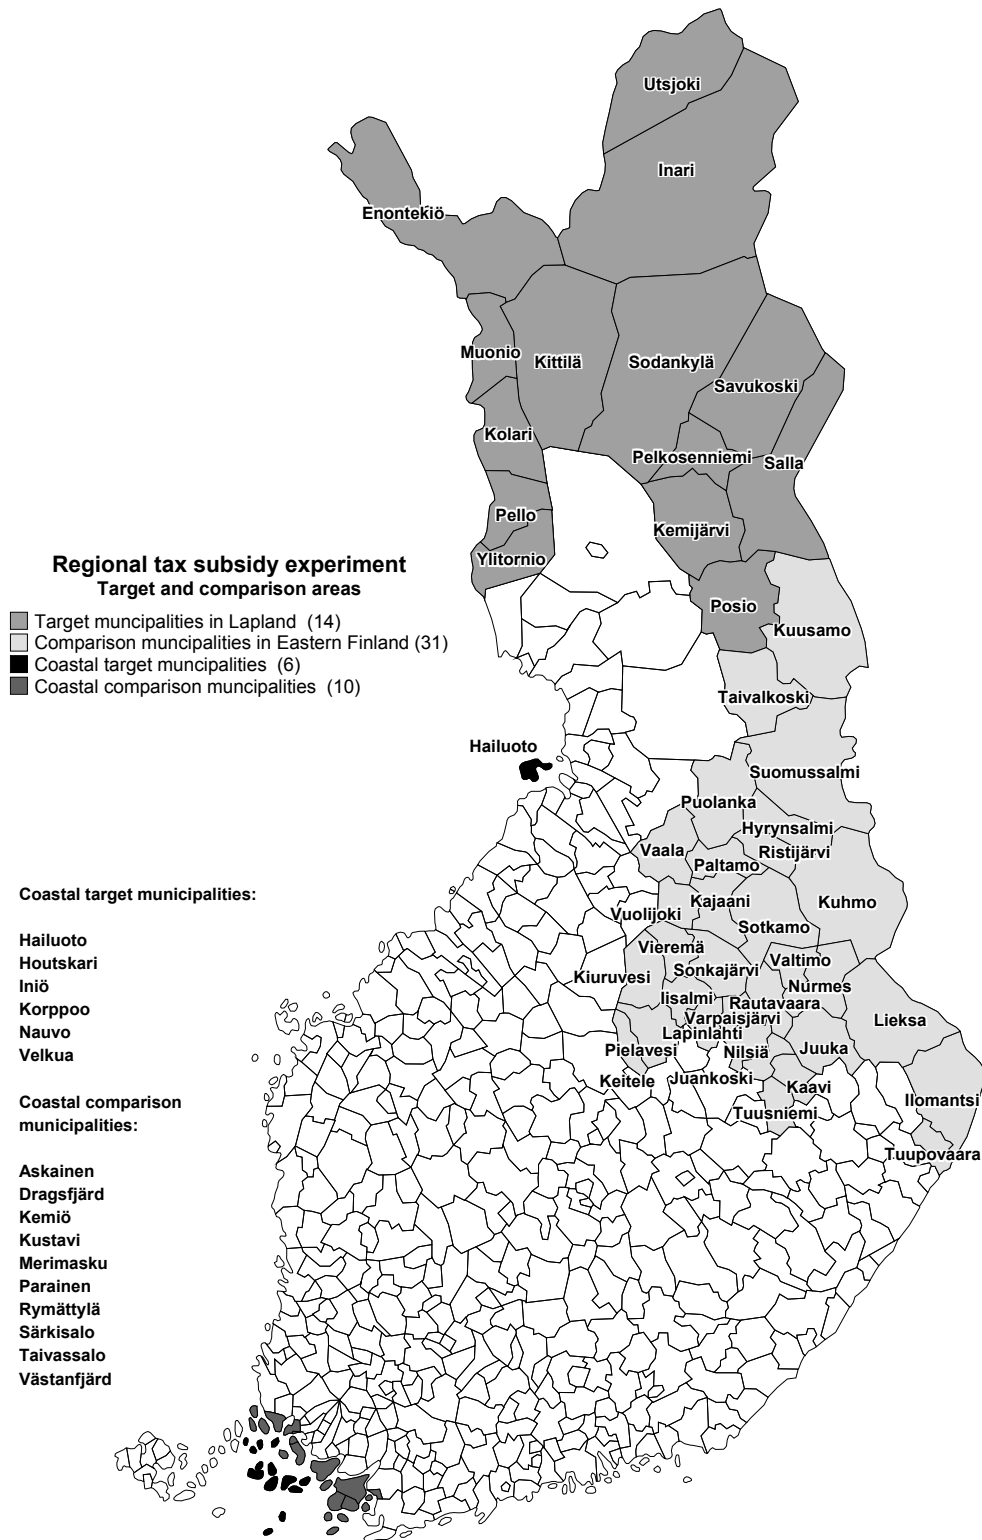


Table 5.2 Comparison of target and control regions


	Regions in Lapland		All Finland
	Target	Comparison	
Population			
Total population	64,979	238,325	5,194,901
Population density 1)	0.84	4.72	17.06
Degree of urbanization 2)	53.28	61.30	83.30
Percent Swedish	0.18	0.07	5.60
Percent Pensioners	27.33	28.57	21.87
Dependency ratio	1.97	1.96	1.30
Secondary education, % 3)	37.85	37.72	36.10
University level education, %	14.85	15.53	23.30
Employment			
Employment rate, %	52.33	53.18	64.16
Unemployment rate, %	23.56	21.27	12.34
Municipal employees, %	22.95	20.56	14.12
Agriculture and fishing, %	11.74	13.39	4.68
Manufacturing, %	8.90	15.96	19.38
Hotels and restaurants, %	6.46	2.90	3.05
Trade, %	9.35	9.39	12.01
Municipal finance			
State grants, € / person	1782	1498	706
Tax revenue, € / person	2085	2022	2715

Notes: 1) inhabitants / km², 2) Indicates the proportion of population living in built-up areas (%), 3) Persons aged 15 or over who have a degree from a senior secondary school, vocational or professional education institution, or from a university.

In addition to having similar population structure and similar industry composition, the aggregate economic development in the target and the comparison regions has been remarkably similar before the experiment. For example, the unemployment rate has a very similar downward trend in the target and comparison regions (Figure 5.2).

The comparison of unemployment rates in Figure 5.2 also indicates that the payroll tax exemption did not have a major impact on unemployment – there is no clear difference between the target and the comparison regions after the beginning of the experiment in January 2003. It would also be interesting to compare changes in employment between the regions, but the available data sources offer limited possibilities for doing that in a reliable way. Since the target region represents only some 1.3% of the Finnish population, the sample sizes in national surveys, such as the Labour Force Survey, become dismally small. Eventually, the problem can be solved by computing regional employment changes based on register data, but currently only data up to 2003 are available.

Figure 5.2 Unemployment rates in the target and the comparison regions



Notes: Employment weighted average of the municipality level unemployment rates reported by the Ministry of Labour. These unemployment rates are calculated by dividing the number of unemployed job seekers in the unemployment register by the number of people in the labour force calculated from administrative data in the end of year $t - 2$.

If the control area is truly similar to the target area, the development in the control area can be used as a valid counterfactual estimate of what would have happened in the target area in the absence of the payroll tax reduction. Careful selection of the comparison region is a necessary pre-condition for the validity of this assumption. While focusing on the employment changes “differences away” pre-existing differences between the target and control regions, it is still possible that the target and the control regions experience different shocks or display different pre-existing trends in employment or wages. In particular, a different industrial structure may lead to different timing of the business cycle in the control and the target regions.

To further enhance the comparability of the target and the comparison regions we matched each firm from the target region with a similar firm or firms from the comparison region. We first split the data into seven main industries using the industry classification in the Labour Force Survey and then applied matching methods to create treatment and comparison groups within these industry classes.

In practice, we estimated logit-models within each industry explaining whether the firm was located in the target region. The explanatory variables were the payroll tax bracket, (log)number of employees, (log)total earnings of the

employees, (log)total sales of the firm (all measured in 1999, 2000, and 2001) and a set of three-digit industry codes. The logit-estimates were then used for calculating for each firm the predicted probability of being located in the target region, i.e. the propensity score.

Each target region firm is then matched with its nearest neighbour (or neighbours) from the comparison region. We used a genetic matching method (Diamond and Sekhon 2005) that uses both the covariates and the propensity score to create matched samples. The genetic matching procedure starts with a weighting scheme identical to Mahalanobis distance. The weight matrix is then iteratively changed using an evolutionary search algorithm (Sekhon and Mebane 1998, Mebane and Sekhon 1998) until no further improvement in match quality is attained (see Diamond and Sekhon for details on match quality criteria). As demonstrated in the appendix of our working paper (Korkeamäki and Uusitalo 2008), this method yields a better match quality with respect to almost all matching variables than simple propensity score matching.

In this evaluation, we will follow the effects during the first two years of the programme. We account for potential anticipatory effects by creating matched samples based on data from the end of 2001, before any information on the programme was made public. To minimize the temptation to re-define the control group ex-post, we fixed the design and published the setup before any data on employment effects became available in January 2004. (Korkeamäki and Uusitalo 2004) The effects of the payroll tax exemption were then evaluated in a transparent way by simply comparing the changes in wages and employment in the treatment and control firms after January 2003. Our last observation date is December 2004. The last year of the experiment is left out because of the changes in the comparison area; the firms in ten municipalities in the Kainuu County that belong to our comparison area became eligible for a similar payroll tax exemption in 2005 as a part of a regional self-government experiment. This new experiment may contaminate the results of the original experiment but it should not be a major issue up to the end of 2004, because adding payroll tax cut to the Kainuu regional self-government experiment was a last-minute change in legislation that was announced only in December 2004.

5.5 Data

We created the matched sample of target and control firms based on the data from the Register of Enterprises and Establishments by Statistics Finland. This register includes data on sales, wage bill, and (imputed) employment of each plant in Finland. Each plant can also be located to a certain municipality. There were 2809 firms in the target area and 7544 firms in the control area. We restricted the sample to the private sector firms that had a positive turnover, paid at least some wages and employed at least one worker in 2001. We also required that the firm has only one plant, so that its location and hence the eligibility for

the tax exemption can be determined accurately. We found 1592 such firms in the target area and 4265 firms in the control area⁶⁹.

The main disadvantage with the establishment register data is that the number of employees in the firm is imputed based on the wage bill, composition of employment and average wages for various employee groups. It is not clear whether the changes in these imputed numbers capture the changes in employment, changes in wages, or perhaps changes in the imputation procedure. Fortunately, comprehensive data on the employment and earnings outcomes was available from the Finnish Tax Administration. The data are based on employer's annual notification that all employers are required to submit to the local tax office. The annual notification includes all wages and salaries paid during the calendar year. The payments are itemized by employee, and the summary form contains the number of recipient itemizations. This number equals the number of employees that have received some wages or salaries from the firm during the year. Naturally, the number of itemizations is only a rough measure of the average employment in the firm. On the other hand, the total wage bill that forms the tax base (i.e. the product of hours worked and the average hourly wage excluding payroll taxes) is accurately reported.

The tax data therefore provides a reliable estimate on whether the payroll tax deduction had an impact on total wage bill. If the total wage bill increased due to the experiment, there must have been an effect on either wages or employment. Reliable estimates on the incidence of payroll tax changes require more detailed information on wages and hours. There is no single database where this information could be gathered for all firms. The best available sources of data on wages and hours are the wage statistics of the employer organizations. In Finland, there are two large employer organizations: Confederation of Finnish Industry and Employers (TT) and Employers Federation of the Service Industries (PT)⁷⁰. Most large employers are members of one of these organizations and the data covers about 60 per cent of private sector employment. Both TT and PT wage surveys contain individual data on all workers in all their member firms. Both surveys contain detailed information on monthly or hourly wages and regular weekly hours. In addition, there are a number of background variables on the employees, including sex, tenure, occupation and industry. More detailed

⁶⁹ The reduction of the sample size is mainly due to dropping firms that had no paid employees in 2001. Many of these firms still had positive turnover. As a robustness check, we included these firms in the sample, but this had no real effect on the results. The sample selection process is described with more detail in Korkeamäki and Uusitalo (2008).

⁷⁰ These two employer organizations merged in 2004. We use data up to 2004 when the wage surveys were still conducted separately.

description of the data is presented in the Appendix of our working paper (Korkeamäki and Uusitalo 2008).

5.6 Results

In the following, we first display evidence that matching balances the characteristics of the firms in the target and the control regions. Then we proceed by presenting the results on the employment changes in the target and comparison regions. We conclude this section with the analysis on wage effects.

5.6.1 Covariate balancing

Table 5.3 reports the means of the variables used in matching separately in the target and comparison regions, and in the matched treatment and control groups. In the rightmost column, we also report the national averages of the same variables. According to the table, the differences between the firms in the target and control regions are rather small to begin with and matching removes most of the remaining differences. A comparison between the treatment and the comparison regions and the national average reveals that both regions differ from the national average and that our comparison region is substantially more similar to the treatment region than to the whole country.

Given the similarity of the target and control regions, there are few strong predictors in the logit-model that is used to explain whether the firm is located in the target region. This is also reflected in the distribution of the propensity score that is rather similar in the target and comparison regions (Figure 5.3). This also implies that finding a region of common support is not a major problem – a large fraction of firms in both regions has an estimated propensity score between 0.1 and 0.5.

As a final check on the comparability between the treatment and the matched control group, we examine the pre-experiment trends in some key variables. Figure 5.4 presents this comparison for the aggregate wage sum. It appears that the firms in the comparison region were larger in the beginning of the period (1996) and have experienced somewhat more rapid growth during the last years of the 1990s. However, the growth in the matched control firms has been very similar to the treatment firms. Note that we use only data from years 1999–2001 in matching and creating the control group, so the similarity in the growth rates before this period is not “forced” into the data, but reflects the similarity between the treatment and the matched control firms. Similar analyses of the long-term trends in mean firm size and aggregate employment did not detect major differences either.

Table 5.3 Covariate balancing

Means, all variables in log's	Target firms	Matched targets	Matched controls	Control region	National average
Employment 2001, SF	1.06	1.06	1.06	1.12	1.26
Employment 2000, SF	1.00	1.00	1.01	1.06	1.27
Employment 1999, SF	0.92	0.92	0.93	0.97	1.25
Employment 2001, TA	1.66	1.66	1.63	1.74	n.a.
Employment 2000, TA	1.49	1.50	1.50	1.58	n.a.
Wage sum 2001, TA	9.65	9.64	9.67	9.74	10.16
Wage sum 2000, TA	8.62	8.64	8.72	8.71	10.13
Wage sum 1999, SF	7.57	7.58	7.60	7.59	10.05
Turnover 2001	11.10	11.12	11.13	10.95	12.16
Turnover 2000	10.36	10.37	10.39	10.08	12.11
Turnover 1999	9.50	9.50	9.57	9.23	12.04
Industry distribution of firms (percent of firms)					
Manufacturing	13.69	13.46	13.46	15.80	11.47
Construction	13.63	13.84	13.84	15.03	13.16
Trade	20.92	21.11	21.11	21.55	16.62
Hotels and restaurants	12.44	12.37	12.37	7.67	4.48
Transport	12.19	12.05	12.05	9.31	10.35
Business services	13.25	13.46	13.46	14.11	20.71
Other services	13.88	13.71	13.71	16.53	22.28
National pension insurance contribution rate					
I (2.95 %)	96.48	96.81	96.81	97.07	92.79
II (5.15 %)	1.01	1.02	1.02	0.98	2.87
III (6.05 %)	2.51	2.17	2.17	1.95	4.34
<i>N</i> Firms 2001	1592	1430	1430	4265	136,434
<i>N</i> Employees 2001, TA	12,318	11,034	10,190	39,111	1,318,654

Notes: For the ease of comparison, we calculated the control group mean displayed in the table using nearest neighbour matching. In the table the industry distribution is reported at a one-digit level. In the actual matching procedure, we use a more detailed industry classification adding 116 three-digit industry codes to the logit-models. The national averages are calculated from the firm register of Statistics Finland for firms with positive employment, wage sum and turnover. SF = Employment figure supplied by Statistics Finland, estimated person-years. TA = Employment figure supplied by Tax Administration.

Figure 5.3 *Estimated propensity score densities*

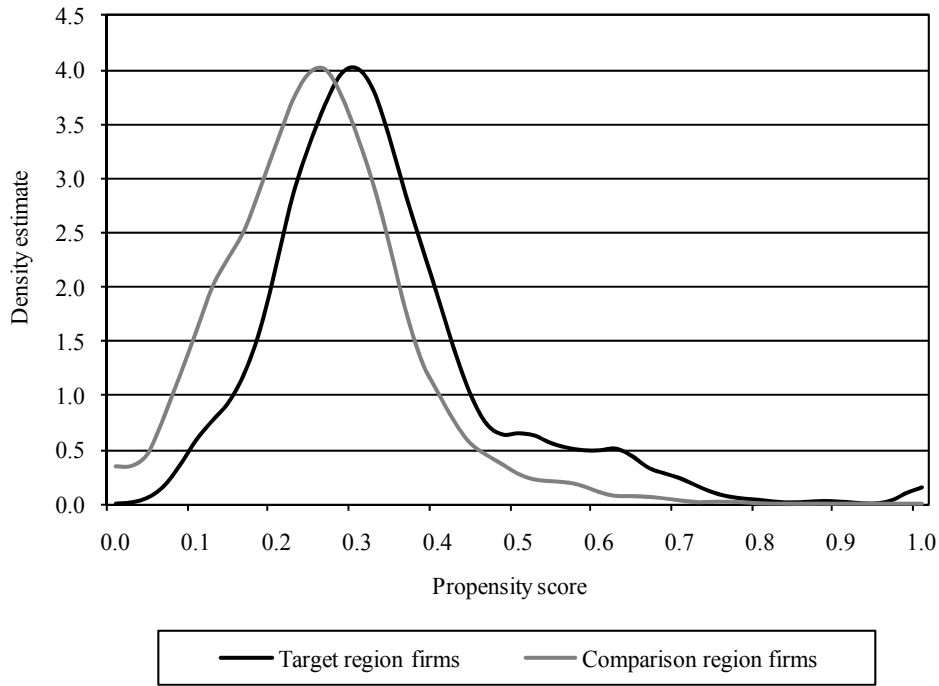



Figure 5.4 *Development of the aggregate wage bill before the experiment*



Notes: Single plant firms that existed in the end of 2001. Comparison region figures are weighted to correspond to the number of firms in the target region at 2001, i. e. weight = number of target firms / number of comparison firms.

5.6.2 Employment and wage sum responses to the regional payroll tax experiment

The main purpose of the payroll tax exemption was to promote employment in the target region. Therefore, also our main outcome measure is the change in the absolute number of employees in a firm. We prefer absolute changes to relative changes – this way employment changes add up to the total effect of the experiment and no weighting is necessary. However, the qualitative results were similar when we use changes in log employment as an outcome measure and weighted the estimates by firm size in 2001.

To reduce noise in employment numbers, we exclude the workers who receive only ancillary income from the firm and concentrate on the employees in their principal employment. Even this measure is naturally imperfect because it does not capture the variation in working hours.

Our main findings on employment effects are reported in Table . The first two columns report the average change in employment in the treatment and control groups. The third column labelled “Treated – Controls” is our estimate for the programme effect. In each case we first report annual changes. In the lower section of the tables, we also calculated two-year changes just before experiment 2000–2002 and after the start of the experiment 2002–2004.

The first observation from Table 5.4 is strong employment growth before 2001 and a strong decrease after 2001 that occurs in both the treatment and the control groups. For example, between 2000 and 2001, employment grew by 0.57 persons in an average treatment group firm. Between 2001 and employment decreased on average by 0.37 persons in the same firms. This pattern is largely due to the entry and exit of firms. Our sample consists of firms that existed in the end of 2001. The firms that exit before the end of 2001, or enter after 2001, are not included in data. On the other hand, exits after 2001 contribute to the average growth rate with large negative changes, and firms that enter before 2001 with large positive changes.

A more important observation from Table 5.4 is that employment growth has been rather similar in treatment and control groups. The differences in growth rates reported in the third column are in most cases smaller than the standard error of the estimate, and in no case anywhere close to being statistically significant. According to these results, the payroll tax experiment has not had a significant effect on employment in the target region. In addition to statistical significance, it is interesting to assess the economic significance of the point estimates. According to Table 5.4 the two-year change in employment after the start of the experiment (2002–2004) was on average 0.103 persons larger in the treatment group. Given that there are 1430 firms in the treatment group, the total

employment effect of the tax cut amounts to 147 new jobs or 1.3 per cent increase in employment.

Table 5.4 Effect of payroll tax cut on employment

	Treated	Controls	Treated – Controls	Std. Error
<u>Change in average number of employees</u>				
2000–2001	0.565	0.550	0.016	0.124
2001–2002	-0.372	-0.289	-0.083	0.131
2002–2003	0.003	0.017	-0.014	0.160
2003–2004	0.204	0.092	0.111	0.160
2000–2002	0.200	0.261	-0.062	0.159
2002–2004	0.218	0.115	0.103	0.219

Notes: The estimates in Table 5.4, Table 5.5 and Table 5.6 are (our favoured) five nearest neighbours estimates, estimated using GenMatch procedure as described in Section 4. Standard errors are robust to heteroskedastic treatment effect.

As noted before, the tax data is not ideal for measuring the changes in employment. On the other hand, any changes in the wage bill (average hourly wage \times sum of hours) that form the tax base should be accurately reported. In Table 5.5 we calculate the effect of the payroll tax cut on the wage bill in the target and control firms. Now the estimates have mostly the “right” sign indicating stronger wage bill growth in the treatment group after the start of the experiment in 2003. The wage bill increase was 1125 euros larger in the treatment group in the first year after the experiment. There was a slight difference also in the second year so that the two-year increase in wage bill was 1728 euros (about 2.7 per cent of an average target area firm wage bill) larger in the treated firms. In addition, these estimates are far from being statistically significant.

Table 5.5 Effect of payroll tax cut on wage bill

	Treated	Controls	Treated – Controls	Std. Error
<u>Average change in wage bill, €</u>				
2000–2001	5328	4689	639	850
2001–2002	1026	1661	-635	945
2002–2003	2263	1137	1125	1597
2003–2004	1666	1063	603	1222
2000–2002	6354	6350	4	1276
2002–2004	3929	2201	1728	2142

Notes: The estimates in Table 5.4, Table 5.5 and Table 5.6 are (our favoured) five nearest neighbours estimates, estimated using GenMatch procedure as described in Section 4. Standard errors are robust to heteroskedastic treatment effect.

5.6.3 The effects by firm type

One could argue that the effect of payroll tax cut might differ across firms. For example, firms paying below average wages may be more responsive to wage costs if the own price demand elasticity of low-skill workers is higher than that of high-skill workers. There could also be different effects in the small and large firms. At least the effect is likely to be smaller in the largest firms that paid more than the deductible maximum of 30 000 euros in payroll taxes. For these firms the payroll tax cut is a lump-sum reduction in taxes and marginal changes in employment should not be affected by the tax rate. Finally, the size of the payroll tax cut depends on the pre-experiment tax-bracket and one might expect larger effects in the firms that face larger payroll tax reductions.

To examine these issues we first split the sample into quartiles defined according to the average wage in the firm and calculated the effects separately in each quartile. In addition, we calculated the effects of the payroll tax cut separately for the firms that paid less than 25 000 euros in payroll taxes in 2001 and that hence were well below the maximum tax deduction. Finally, we calculated the effects for the firms that were in the lowest payroll tax bracket. (The number of firms in the higher brackets was too small for meaningful calculations).

Table 5.6 Effect of payroll tax cut by firm type

	Treatment – control difference in			
	Empl. change 2000–02	Empl. change 2002–04	Wage bill change 2000–02	Wage bill change 2002–04
Full sample	-0.062 (0.159)	0.103 (0.219)	4 (1276)	1728 (2142)
By wage quartile				
1 st (lowest)	-0.027 (0.139)	-0.210 (0.133)	-2086 (604)	-281 (612)
2 nd	-0.119 (0.173)	0.724 (0.204)	218 (1122)	2666 (1280)
3 rd	-0.264 (0.273)	0.457 (0.272)	1124 (2083)	9217 (3226)
4 th (highest)	-0.245 (0.177)	0.000 (0.296)	55 (1999)	356 (4826)
Firms in lowest payroll tax bracket	-0.122 (0.121)	0.160 (0.159)	66 (881)	1264 (1274)
Firms paying less than 25 000 € in payroll taxes	-0.122 (0.134)	0.026 (0.175)	559 (138)	-455 (1775)

Table 5.6 reports the results of these experiments. No clear patterns appear. The effect of the payroll tax cut on employment seems to be the highest in the second wage quartile. The effect seems also to be higher than full sample average in the small firms that are in the lowest payroll tax bracket and in the firms that pay less

that 25 000 euros in payroll taxes. The effects on the wage bill change appear rather similar, though now the largest positive effects appear in the third wage quartile. Due to large standard errors associated with all sub-sample estimates, not much can be concluded from these numbers.

5.6.4 The effect on wages

To have a closer look at the incidence of the payroll tax cut we examined its effect on hourly wages. As noted before wage data is available only for the subset of (large) firms that belong to one of the two employer organizations. These two organizations have slightly different surveys and different wage concepts. The manufacturing sector data is also divided into the white-collar and the blue-collar worker files according to whether the employees receive monthly salaries or hourly wages. To avoid the need of ad hoc adjustments for different measurements, we also report the estimates separately. For the service sector workers and for blue-collar workers in manufacturing we have data for the period from 2000 to 2004, for the white-collar workers in manufacturing only for 2001–2004.

While the firm is a natural unit of observation when measuring changes in employment, it is more straightforward to use individual wages to estimate average wage growth. Our wage equation estimates are reported in Table 5.7. In each case we create a measure that accounts for the variation in working hours. For the workers that receive monthly salaries we divide monthly salary by usual hours. For workers that are paid by hour we divide total wages during the last quarter of the year by total hours during the same period. We estimate the wage equations using all wage components (including various bonuses). To account for unobserved individual-level variation in wages we use data for the employees who appear in the data in the two consecutive years and use the change in real log wage as a dependent variable.

All wage equations include the usual control variables: age, education and gender. We also include an indicator for supervisory or trainee status when available, and add a full set of two-digit occupational dummies in the wage equations. The equations include year fixed-effects as well as a fixed-effect for being located in the target region. The effect of the payroll tax cut is identified from the interaction between year 2003 and target region indicators. The coefficient of this interaction can be interpreted as the difference in wage growth rate between the employees in the target and control regions due to the start of the experiment. Note that the interaction between year 2004 and the target region should be zero unless wage adjustments involve long lags since there were no changes in payroll taxes between 2003 and 2004.

Table 5.7 Wage effects

	Service sector		Manufacturing, salaried		Manufacturing, blue-collar	
Year 2002	-0.005 (0.006)	-0.006 (0.005)	n.a.	n.a.	0.008 (0.025)	-0.002 (0.023)
Year 2003	-0.009 (0.004)	-0.010 (0.004)	-0.012 (0.013)	-0.012 (0.012)	-0.040 (0.024)	-0.046 (0.024)
Year 2004	-0.009 (0.005)	-0.011 (0.004)	-0.012 (0.015)	-0.011 (0.014)	0.013 (0.022)	0.007 (0.019)
Target region	-0.003 (0.007)	-0.007 (0.004)	-0.011 (0.010)	-0.004 (0.004)	-0.000 (0.020)	-0.041 (0.013)
Target region × 2002	-0.001 (0.010)	—	n.a.	n.a.	-0.073 (0.029)	—
Target region × 2003 (Treatment effect)	0.016 (0.010)	0.020 (0.010)	0.015 (0.011)	0.008 (0.006)	0.025 (0.028)	0.066 (0.028)
Target region × 2004	-0.010 (0.010)	—	0.014 (0.015)	—	-0.047 (0.029)	—
R^2	0.028	0.027	0.070	0.070	0.042	0.040
N obs.	9972	9972	2493	2493	9721	9721
N indiv. t-region	746	746	108	108	408	408
N indiv. c-region	3134	3134	1028	1028	3133	3133
N firms t-region	81	81	8	8	11	11
N firms c-region	255	255	39	39	45	45

Notes: The dependent variable in all regressions is the change in log hourly wages including overtime, benefits (taxable value), and provision payments. All equations include gender, age, age squared, dummies for occupations (67 for manufacturing 41 for service sector). Service sector regression has additional dummy variables for trainees, supervisors, and managers. Both manufacturing sector regressions have controls for education level. Robust standard errors are calculated taking into account clustering by firm. From manufacturing sector we have excluded one large target region firm that shut down during the observation period. n.a. = not available due to lack of data.

According to results in column 1 of Table 5.7, service sector wage growth seems to have been very similar in the target and control regions before the experiment started. In 2003, when the payroll taxes were cut, wages grew 1.6 per cent faster in the target region though the estimate is not statistically different from zero. The point estimates also suggest that wage growth was slightly slower in the target region in 2004 but also these estimates are insignificant.

The specification including all interaction terms reported in column 1 effectively compares the differences in wage growth between the reform year 2003 and the base year 2001. In column 2, we report results from a specification that restricts all the other interactions except the interaction between 2003 and the target group to zero, effectively comparing wage changes in the reform year to all other years. The estimate is now slightly higher and statistically significant indicating that the tax exemption led to 2 per cent faster wage growth among the employees in the

treatment group. This result is robust to small changes in the model specification such as restricting the sample to occupations that are present in both target and control groups, measuring occupations at three-digit level or excluding bonuses from the wage measure.

In the manufacturing sector the number of target group firms is smaller and the results are sensitive to whether one large firm that closed down during the period is included in the data or not. The estimates are also generally less robust to small changes in specification. In fact, the largest wage changes in manufacturing – such as the relative decline in wages in 2002 – seem to be unrelated to the reform.

Above, we estimated all wage equations at the individual level. This may be problematic since changes in the large firms have a large weight in the estimates. If there are firm-specific shocks, the results may be driven by the shocks that occur in some large firms. To reduce the weight of these large firms we experimented with re-weighting the data so that each firm gets the same weight. Except for the blue-collar workers (where one large firm dominated the results), this re-weighting had only a minor effect. In particular, the result that the service sector wages grew slightly faster in the target region was robust to re-weighting.

5.7 Concluding comments

Well designed policy experiments may provide valuable information for policymakers on the effects of taxation on wages and employment. In an ideal case, estimates based on regional experiments are more reliable than estimates based on cross-country comparisons or time-series data. Estimates from these experiments could then be used for cost-benefit analysis and as a basis for future tax policy. The main problem in small-scale experiments tends to be the small number of observations. Measurements of the employment changes in firms are noisy and pinning down the effects of reasonably small changes in payroll taxes would require a large experiment.

The Finnish payroll tax experiment reduced payroll taxes by 4.1 per cent, on average. If the estimates for the sub-sample of firms for which wage data is available can be generalized to all firms, about half of the effect of the payroll tax reduction on labour costs was offset by faster wage growth in the firms that were eligible for the payroll tax cut. The remaining two per cent decrease in labour costs did not have a significant effect on employment, but the estimates are not very precise due to the small sample size. Still, our point estimates of tax incidence are somewhat different from earlier results by Gruber (1994, 1997) and Johansen and Klette (1998), according to which reductions in payroll taxes are almost entirely shifted to wages. These studies imply that labour supply is less elastic than labour demand, while according to our estimates the demand and supply elasticities are roughly equal. According to our point estimates, the tax cut

increased employment by 1.3 per cent, indicating that labour demand elasticity is around 0.6, well within the range of earlier estimates. Unfortunately, the confidence bands around this estimate are too wide to give much guidance for future tax policy.

The fact that the cut in payroll taxes was targeted at narrowly defined regions, and the temporary nature of the tax cut, naturally limits the extent to which the results can be generalized to the potential effects of permanently reducing payroll taxes in the whole country. First, the payroll tax experiment was financed by increasing payroll taxes in the rest of the country. In a national scheme, the budgetary cost would need to be financed by raising other taxes. Second, a regional experiment may have substitution effects if firms reallocate labour to the target region from the rest of the country. This might be beneficial in the sense that part of the reasons for the regional payroll tax cut was to boost employment in disadvantaged regions. However, this limits the usefulness of the results from the experiment in predicting the effects of a national programme. Third, the incidence of the tax cut may also be different in a regional programme since union contracts are negotiated at the national level. Any nationwide changes in payroll taxes may have an impact on the outcome of these negotiations, while a regional programme that only affects a small share of employers has little weight in national bargaining. Finally, a temporary programme is likely to create smaller employment effects than a permanent reduction in payroll taxes. Three years may not be a sufficiently long period for firms to adjust their labour demand to a relatively small change in labour costs.

References

- Anderson, P. M. – Meyer, B. D. (1997): The effects of firm specific taxes and government mandates with an application to the U.S. unemployment insurance program. *Journal of Public Economics* 65(2), 119–145.
- Bennmarker, H. – Mellander, E. – Öckert, B. (2009): Do regional payroll tax reductions boost employment? *Labour Economics* 16, pp. 480–489.
- Bohm, P. – Lind, H. (1984): Sysselsättningseffekter av sänkt arbetsgivaravgift i Norrbotten/Svappavaara: Metodbeskrivning, SIND PM 1984:2, Stockholm.
- Bohm, P. – Lind, H. (1993): Policy evaluation quality. *Regional Science and Urban Economics* 23(1), 51–65.
- Diamond, A. – Sekhon, J. S. (2005): Genetic matching for estimating causal effects: A general multivariate matching method for achieving balance in observational studies, Working Paper, <http://sekhon.polisci.berkeley.edu/papers/GenMatch.pdf>

- Gruber, J. (1994): The incidence of mandated maternity benefits. *American Economic Review* 84(3), 622–641.
- Gruber, J. (1997): Tax incidence of payroll taxation: evidence from Chile. *Journal of Labor Economics* 15(3), 72–101.
- Johansen F. – Klette, T. J. (1998): Wage and employment effects of payroll taxes and investment subsidies. Department of Economics, University of Oslo, memorandum 27/1998.
- Hamermesh, D. (1993): Labor Demand. Princeton University Press.
- KELA (2002): Statistical yearbook of the Social Insurance Institution, Finland, a publication by the Social Insurance Institution T1:38.
- Mebane, W. R. – Sekhon, J. S. (1998): GENetic Optimization Using Derivatives (GENOUD). Software Package. <http://sekhon.polisci.berkeley.edu/rgenoud/>.
- Murphy, K. J. (2007): The impact of unemployment insurance taxes on wages. *Labour Economics* 14(3), 457–484.
- Korkeamäki, O. – Uusitalo, R. (2004): Työnantajan sosiaaliturvamaksusta vapauttamisen alueellisen kokeilun työllisyysvaikutukset, Väiliraportti I. Discussion Papers 194, Labour Institute for Economic Research, Helsinki.
- Korkeamäki, O. – Uusitalo, R. (2008): Employment and wage effects of payroll tax cut – evidence from a regional experiment. VATT Working Papers 443, Government Institute for Economic Research.
- Rosenbaum, P. R. – Rubin, D. B. (1985): Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. *The American Statistician* 39(1), 33–38.
- Sekhon, J. S. – Mebane, W. R. (1998): Genetic Optimization Using Derivatives: Theory and Application to Nonlinear Models. *Political Analysis* 7, 189–203.
- Uusitalo, R. (2005): Do centralized bargains lead to wage moderation? Time-series evidence from Finland. In Piekkola, H. – Snellman, K. (Eds.): Collective bargaining and wage formation. Physica-Verlag, Heidelberg.
- Uusitalo, R. – Vartiainen, J. (2007): Finland: firm factors in wages and wage changes. In Lazear, E. – Shaw, K. (Eds.): Wage structure, raises and mobility: international comparisons of the structure of wages within and across firms. University of Chicago Press (in press).

6. The Finnish payroll tax cut experiment revisited, or where did the money go?⁷¹

Abstract

In this paper I evaluate the effects of a regional experiment that reduced payroll taxes by 3–6 percentage points of the firms' wage sum in northern and eastern Finland. I estimate the effect of the payroll tax reduction on firms' employment levels, wage sum and profits, and on workers' hourly pay and monthly hours worked, by comparing the changes in employment and wages before and after the start of the experiment to a comparison region. My results indicate that the reduction in payroll taxes did not lead to any unequivocal aggregate effects in the target region.

Key words: payroll tax, labour demand, tax incidence

JEL classification: J18, J23, J38, J58, J65, J68

6.1 Introduction and background

The Finnish payroll tax experiment that started in 2003 was originally limited to a three-year period and the evaluation is presented in the previous chapter. Extension of the experiment until 2012 and the enlargement of the target region and the target population of firms to cover almost twice the area and more than twice the number of firms warrant a further investigation. In this chapter I take different, more straightforward tack on methods, following Benmarker, Mellander and Öckert (2009), and include also the firms' profits into the set of response variables.

There is a rather strong consensus regarding the labour market effects of payroll taxes. The textbook model states that a reduction in payroll taxes lowers wage costs and hence boosts the demand for labour. Its effect on employment then depends on the incidence of the tax. If the tax cut leads to higher wages that entirely offset the reduction in taxes, the tax cut will have no effect on employment, and if the labour supply is fully elastic, then the tax cut will result in higher employment. The general finding of recent empirical research is that

⁷¹ This is a somewhat shortened version of the similarly titled working paper (Korkeamäki, 2011). This research was financed by the Ministry of Employment and the Economy. I am grateful to Roope Uusitalo and Kari Hämäläinen for their comments and suggestions, which considerably improved the paper.

changes in payroll taxes are partly shifted to wages, with little effect on employment⁷².

Textbooks, however, say very little on the effect of payroll taxes on firm profitability. In the neoclassical family of labour market models, the zero profit constraint seems to void the question altogether. In recent years, empirical observations have dented these theories to some extent⁷³. In their book on a closely related subject, minimum wages⁷⁴, Card and Krueger (1995) consider the possible mechanisms by which changes in minimum wages could affect firms' profits⁷⁵. Lacking suitable micro data for direct measurement, they use changes in firms' stock market valuations as an indicator of changes in profits. The changes in minimum wage legislation (or announced changes) are used as instruments to identify the effect of minimum wages on stock prices, and the implied change in profits is calculated. Card and Krueger find tentative evidence that announced rises in minimum wages induce investors to adjust their valuation of firms downward.

The first study on the direct effect of minimum wages on firm profitability is Draca, Machin and Van Reenen (2008). They use the introduction of a national minimum wage to the UK labour market in 1999 as a quasi-experiment to identify the effect of a rise in minimum wages on profits. The motivation for their study is that in the UK case there was little impact on employment (Machin, Manning and Rahman 2003 and Stewart 2004) and also little evidence that firms were able to pass on higher costs to consumers by increasing prices (exceptions here are Aaronson 2001 and Aaronson and French 2007). Draca et al. find a significant reduction in profits and a rise in labour costs owing to the introduction of a national minimum wage scheme, but neither employment nor productivity changed. They also report that in the longer run the labour cost hike did not seem to force the affected firms out of business.

While there have not been any abrupt changes in wage schemes, the ongoing experiment in payroll taxes⁷⁶ could be used as an instrument to estimate a wage

⁷² See Benmarker *et al.* for a short review of previous studies.

⁷³ There is a quotation attributed to Paul A. Samuelson that "In economics it takes a theory to kill a theory; facts can only dent a theorist's hide."

⁷⁴ In the part of the wage distribution where minimum wage rules are binding, the effect of minimum wages can be considered to be similar to a payroll tax hike. The main difference is that the uneven incidence can cause substitution away from low-wage labour towards both capital and higher-wage labour.

⁷⁵ The focus of the book is on the employment effects of minimum wages, but there is a chapter on how much profits change. Unlike the payroll tax case, standard economic theory unambiguously implies that wage floors have a negative impact on employment (Borjas 2005, Brown 1999). Empirical evidence is considerably more mixed; see the comprehensive review by Neumark and Wascher (2007).

⁷⁶ See the previous chapter for a description of the original experiment.

cost effect on firms' profits. The payroll tax exemption was planned to last for three years, from 1 January 2003 to 31 December 2005. Already in May 2003, the government had decided to start a regional self-government experiment in Kainuu, eastern Finland, beginning from 2005. That experiment contained a similar provision for lowered payroll taxes as the Lapland experiment and hence the payroll tax experiment was expanded and extended to the end of 2009. The experiment has since been extended further until the end of 2012.

To sum up the current situation and motivate the need to assess whether payroll tax cuts have had an effect on firm profits, I draw the following conclusions. 1) According to our previous research of the first two years of the payroll tax experiment, the cut in northern Finland did not seem to have any immediate employment effects (Korkeamäki and Uusitalo 2009). This finding is supported by evidence from other Nordic labour markets. 2) There was some indication of rising wages, but not 1:1 with respect to the tax break – this is a finding that has also been made in Sweden and Norway. From 1) and 2) and supported by the UK case of a change in minimum wages, it seems likely that changes in payroll taxes could have an effect on firm profitability. Models of incomplete competition from the IO literature (Aaronson and French 2007) and matching models from the labour market side (e.g. Flinn 2006) can accommodate these profit effects, but their size remains an empirical question.

For this study, I have twice the number of firms in the treatment group compared to the earlier research, better data and more years of observations. However, I still do not find any effects on employment, wage sum or profits. The wage sum and profits measured in euro terms grew faster in the target region of the experiment, whereas the employment gains were negative, but none of the effects are statistically significant. The additional information available here does make the previous results concerning wages suspect, however – there still is a positive and significant wage effect in Lapland, but in Kainuu the effect is negative and significant. Certainly, there might have been a region-specific shock in Kainuu causing the negative effect, but I found no reason to believe that the result for Lapland was trustworthy.

That there was a tax cut is a fact and it can be observed to have lowered the cost of employment. Other results, however, are either non-existent or drowned in the standard errors. It is unfortunate that researchers were not consulted in the design phase of the experiment. The selection of the target region and the size of the tax cut were mainly driven by political feasibility, not by a focus on facilitating reliable and conclusive research.

6.2 The experiment, target and comparison regions and firms

At the turn of the millennium there was an ongoing debate over the relative merits of across-the-board, low-bureaucracy tax cuts and more targeted measures

to promote employment. In March 2002, the Finnish government agreed to a temporary removal of employer contributions to national pension insurance and national health insurance (see Table 6.1) for firms operating in the 20 target municipalities⁷⁷. The programme was designed as an experiment with the stated aim of evaluating the effect of a cut in payroll taxes on employment in the target region. The tax cut was designed to fit within the European Union *de minimis* regulations that govern firm subsidies. Therefore the maximum tax cut is 30,000 euros per year for each firm and the already heavily subsidised industries of agriculture, fishing and transport were excluded from the experiment. The payroll tax exemption was to continue for three years from January 1 2003 to December 31 2005. In December 2005, the government extended the duration of the experiment to the end of 2009. The original regional tax experiment is exhaustively described in Korkeamäki and Uusitalo (2009).

The act on the regional self-government experiment in Kainuu was passed by the Finnish parliament in February 2003 and the experiment started on 1 January 2005. The aim of the self-government experiment is to gain experience of the effects of regional self-government on regional development work, basic services, citizen activity, the relationship between regional and state central government as well as between municipal and state local government. The Kainuu experiment provides the same payroll tax cut as the Lapland experiment but it is no longer motivated in the law as being an experiment nor is there any specific mention of an evaluation of the tax cut. The Kainuu experiment extends the payroll tax cut to public sector employers and this provision was extended to Lapland from the beginning of 2006.

The Kainuu region has nine municipalities with an area nearly equalling that of Belgium, but a population of only 85,000. The target region in Lapland is even larger in area, with a population of 65,000. Both can be described as sparsely populated, high unemployment regions with little manufacturing or other industrial activity. The share employed in agriculture and forestry is much higher and the average level of education much lower than in the rest of the country. The biggest employer is local government.

6.2.1 Finnish payroll taxes

Payroll taxes in Finland consist of employer contributions to the employees' pension scheme, national pension insurance, national health insurance, employment accident insurance, and unemployment insurance. The tax rates for the various components vary across sectors and by firm size, and firms' pension

⁷⁷ The target region for the original experiment was 14 municipalities in Lapland and 6 municipalities on the islands off the south-west coast of Finland.

contributions depend on the characteristics of their employees. The components of the payroll tax and their evolution over the 15 years from 1995 to 2009 are presented in Table 6.1. The largest component – contributions to the employees' pension scheme – has remained stable, while the other components have gradually been lowered after the recession in the early 1990s.

Table 6.1 The components of Finnish payroll taxes, percentage of the wage sum

Date of change	Employees' pension scheme	National pension insurance + national health insurance			Accident insurance	Unemployment insurance		Group life insurance	Total	
		I	II	III		Part of wage bill under € 840,940	Part of wage bill over € 840,940		Low	High
1.1.1995	16.60	4.000	5.600	6.500	1.2	2.00	6.10	0.120	23.920	30.520
1.1.1996	16.80	4.000	5.600	6.500	1.2	1.00	4.00	0.100	23.100	28.600
1.1.1997	16.70	4.000	5.600	6.500	1.4	1.00	4.00	0.090	23.190	28.690
1.1.1998	16.80	4.000	5.600	6.500	1.4	0.90	3.90	0.080	23.180	28.680
1.1.1999	16.80	4.000	5.600	6.500	1.3	0.90	3.85	0.080	23.080	28.530
1.1.2000	16.80	4.000	5.600	6.500	1.2	0.90	3.45	0.090	22.990	28.040
1.7.2000	16.80	3.600	5.600	6.500	1.2	0.90	3.45	0.090	22.590	28.040
1.1.2001	16.60	3.600	5.600	6.500	1.2	0.80	3.10	0.095	22.295	27.495
1.1.2002	16.70	3.600	5.600	6.500	1.1	0.70	2.70	0.095	22.185	27.085
1.3.2002	16.70	2.950	5.150	6.050	1.1	0.70	2.70	0.095	21.535	26.635
1.1.2003	16.80	2.964	5.164	6.064	1.1	0.60	2.45	0.081	21.545	26.495
1.1.2004	16.80	2.964	5.164	6.064	1.1	0.60	2.50	0.080	21.544	26.544
1.1.2005	16.80	2.966	5.166	6.066	1.2	0.70	2.80	0.080	21.746	26.946
1.1.2006	16.70	2.958	5.158	6.058	1.1	0.75	2.95	0.080	21.588	26.888
1.1.2007	16.64	2.951	5.151	6.051	1.1	0.75	2.95	0.080	21.521	26.821
1.1.2008	16.80	2.771	4.971	5.871	1.0	0.70	2.90	0.080	21.351	26.651
1.1.2009	16.80	2.801	5.001	5.901	1.0	0.65	2.70	0.070	21.321	26.471
1.4.2009	16.80	2.000	4.201	5.101	1.0	0.65	2.70	0.070	20.520	25.671

Notes: contribution to employees' pension scheme is the average percentage share. The actual contribution depends on firm size and the characteristics of employees. The cost of accident insurance is also an average.

6.2.2 Target and comparison regions used in the evaluation

In our evaluation of the beginning of the Lapland experiment, the comparison region we chose was in northern Finland in an area with municipalities with similar economic and demographic conditions to those in the original target region. However, the core of our comparison region was Kainuu. Therefore, it became necessary to select a new region to work as a counterfactual for the larger experiment region.

Rather than hand-picking municipalities, I followed Benmarker et al. and used the national firm subsidy rules to find an area where the operating environment

for firms is comparable to the target region. The target region is contained in the two highest subsidy regions for the period 2000–2006 and in the highest category for 2007–2013. In the first period, firms in the Kainuu region and its surroundings to the west and south were eligible for the highest subsidies, with Lapland belonging to the second category. There was, however, a special provision for Lapland that granted firms almost the same investment and other subsidies as for firms in the first category⁷⁸. The subsidy regimes were allocated according to EU rules, where the main factor was the level of NUTS3 region GDP per capita relative to the EU average – regions with less than 75% of the average were eligible for the highest subsidies.

The comparison region is formed of the non-target municipalities of the two highest subsidy regions for the period 2000–2006. I have excluded the largest local administrative centres and university towns (Rovaniemi, Joensuu, Kuopio and Mikkelä) and one highly industrialised region (Kemi-Tornio) as there is nothing comparable in the target region. Figure 6.1 shows the regions on a map. I decided to drop the target region in the archipelago from this evaluation since it would have been hard to find a credible comparison for this very distinct group of municipalities.

Table 6.2 highlights some important similarities and differences between the target and comparison regions and contrasts them with the rest of the country. The figures are from 2001, i.e. before the experiment had begun, but the main features are quite persistent through the whole period under evaluation. First, the part of Lapland that received the tax cut and Kainuu are very sparsely populated. The comparison region has more than four times as many inhabitants per square kilometre. However, the rest of Finland is more than five times as densely populated as the comparison region. Second, the population in both the target and comparison regions is declining and not growing. It is also older and less educated than the rest of the country. Third, the employment rate was markedly lower (and unemployment rate higher) in the target and comparison regions than in other parts of Finland. The share of municipal employees is particularly high in the Lapland and Kainuu regions and the share in the comparison region does not quite match that. The employment share of manufacturing is clearly lower in the target region but the shares of other industries are well aligned. I will look at the industry composition more closely when I describe the firms in the target and comparison regions. Last, if we consider the public finance situation in the region that received the tax cut, we can see that the target and comparison regions are

⁷⁸ The subsidy scheme is quite complex (details in the Aid to Business Act, 1200/2000). To simplify, the highest share of investment subsidies in category I is 40% and in the northern part (Lapland) of category II it is 34% of the total investment.

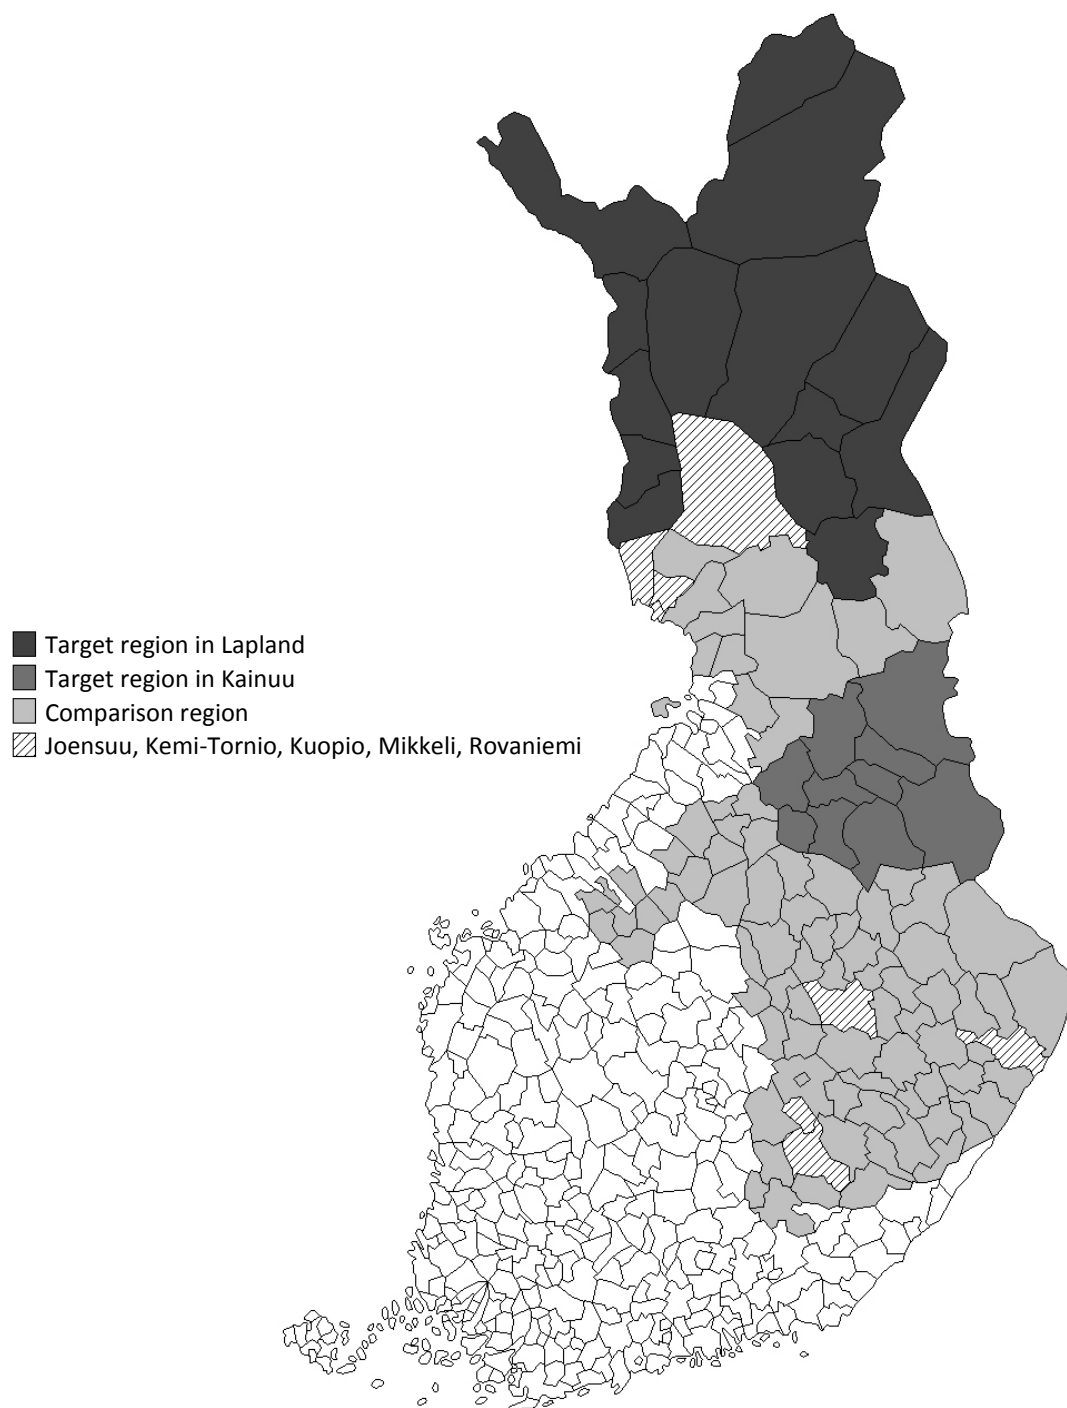
very alike: both are heavily dependent on state grants to finance their public sector.

Table 6.2 Target and comparison regions in 2001

	Target	Control	Rest of Finland
Population			
Total population	153,452	522,418	4,500,946
Population density ¹⁾	1.56	6.54	35.94
Population growth, % / a	-1.80	-0.98	0.48
Percentage pensioners	26.96	27.78	21.01
Dependency ratio	1.94	1.85	1.28
Secondary education, % ²⁾	38.15	37.60	35.68
University level education, %	16.72	15.40	25.14
Employment			
Employment rate, %	52.33	55.98	65.70
Unemployment rate, %	21.05	16.30	11.17
Municipal employees, %	21.98	18.43	13.47
Agriculture, forestry and fishing, %	10.02	13.97	3.61
Manufacturing, %	20.75	25.84	27.00
Trade, %	9.56	9.16	12.33
Municipal finance			
State grants, € / person	1,591	1,399	593
Tax revenue, € / person	2,134	2,007	2,807

Notes: 1) inhabitants / km², 2) Persons aged 15 or over with a degree from an upper secondary school, vocational or professional education institution, or a university. Source: ALTIKA regional statistics database by Statistics Finland.

Figure 6.1 Target and comparison regions



Notes: Lapland and Kainuu form the target region of the tax cut. Joensuu, Kemi-Tornio, Kuopio, Mikkeli and Rovaniemi are removed from the comparison region.

6.2.3 Target and comparison firms

Here I take a short look at the firm population in the tax cut's target and comparison regions. The first observation is that the firms are small – none of the firms that have all of their establishments situated in the combined area of the target and comparison regions has more than 600 employees. Furthermore, none of the firms in the target region has over 300 employees. This leads me to make one common support-type restriction for the comparison group: I drop a few large firms from the comparison group, as it is unclear if they are comparable to any firms in the target area. Other restrictions have to do with EU regulations on firm subsidies (firms in agriculture, fisheries and transport are not eligible for the payroll tax cut) and the technical properties of the firm and establishment data. I use only observations for firms that can be reliably linked between different registers and to all of their establishments for each year they occur in the datasets. In addition, I require that the information from all sources on the key variables is consistent⁷⁹.

The main firm-level response variables in this study are employment, wage sum and operating profit. Almost all the other variables, e.g. various attributes of the firms' workforces and financial position, are more or less endogenous and hence cannot be used as explanatory variables. Were this a matching exercise, however, these and other pre-experiment firm characteristics would be used to first match and then to assess the quality of the matches. Therefore, I gauge the validity of the quasi-experimental setting in a similar manner by comparing the target area and comparison area firm populations. The only comparison variables in the regressions I run are industry and firm age group dummies. Even though the difference-in-differences set-up should remove time-constant firm-specific (and hence region-specific) differences in levels, dissimilarities in industry growth trends should be taken into account if there are differences in the industry distributions between the target and comparison region firm populations. I report these distributions in Table 6.3. In addition to the distributions, I also calculated a normalised difference for each industry share: Imbens and Wooldridge (2009) consider this a good measure to evaluate whether the regression methods are well suited to estimate the treatment effects. Imbens and Rubin (forthcoming) argue that normalised differences exceeding one quarter in absolute value would probably indicate problems. I also calculated a *t*-statistic for each variable. If this is a reasonable thing to do for a set of inter-related dummies might be questioned, but in the case of the industry distributions, it is not of importance if they differ

⁷⁹ Observations are dropped if there is conflicting information on the same variable from different sources. For example, if according to Financial Statements data a firm has three establishments but not all of those are given in the Business Register, or if there are large discrepancies in total wages or turnover from different sources, the observation is removed.

in a statistically significant manner. The main point is to show that the distributions are similar enough that after controlling for industry the comparison between regions is internally valid.

Table 6.3 *Industry distribution of target and comparison region firms in 2001*

	Target	Control	Normalised difference	<i>t</i> -statistic
Mining and quarrying	0.011	0.020	-0.054	-4.39
Food, beverages and tobacco	0.021	0.021	0.002	0.12
Clothes, etc.	0.009	0.010	-0.005	-0.40
Wood, paper, etc.	0.042	0.040	0.006	0.45
Petro-chemical, etc.	0.002	0.004	-0.022	-1.83
Non-metallic mineral products	0.007	0.008	-0.008	-0.65
All metal industries, except ↙ Electronic and optical products	0.037	0.063	-0.085	-6.83
Water and electricity supply	0.011	0.011	-0.002	-0.14
Construction	0.147	0.170	-0.044	-3.41
Trade of gasoline, repair & trade of motor vehicles	0.044	0.054	-0.031	-2.38
Wholesale and retail trade	0.180	0.175	0.008	0.62
Accommodation and restaurants	0.103	0.071	0.079	5.73
Information and communication	0.126	0.100	0.060	4.38
Finance and banking	0.001	0.001	0.007	0.53
Business services	0.129	0.128	0.001	0.07
Other services	0.125	0.116	0.020	1.46

Notes: Normalised difference is the difference in sample means scaled by the root of the sum of the sample variances, i.e. $\Delta_x = \frac{\bar{X}_T - \bar{X}_C}{\sqrt{S_T^2 + S_C^2}}$ and $t = \frac{\bar{X}_T - \bar{X}_C}{\sqrt{S_T^2/N_T + S_C^2/N_C}}$. Subscript *T* refers to the target group and *C* to the comparison group.

According to Table 6.3 there are some statistically significant differences in the industry dummies but the standardised differences are well under the aforementioned 0.25 in absolute value. Table 6.4 reports the pre-experiment values of the dependent variables (and turnover). Here it might be argued that the *t*-statistic is the more interesting measure. If, indeed, there are significant differences (in differences) in the main outcomes immediately before the experiment, a possible point of concern is whether controlling for firm fixed effects is enough to make causal inferences on the effects of the tax cut valid. In Table 6.4 there are no statistically significant differences at the one per cent risk level. The target region firms are somewhat smaller and their growth two years prior to the start of the experiment in Lapland was a little slower than in the comparison region. Even if this difference is not significant, this might call for the use of firm-specific slopes in the regressions.

Table 6.4 Pre-reform comparison of key variables for target and comparison region firms

	Mean target	Mean control	Normalised difference	t-value	N obs. target	N obs. control
Employment ¹⁾						
2001	3.40	3.70	-0.022	-1.56	2,933	8,851
2000	3.40	3.69	-0.022	-1.54	2,894	8,620
1999	3.41	3.65	-0.019	-1.31	2,707	8,117
Employment growth ²⁾						
2000–2001	0.01	0.00	0.002	0.13	2,665	8,006
1999–2000	0.04	0.13	-0.028	-1.78	2,597	7,810
1999–2001	0.05	0.12	-0.019	-1.22	2,453	7,469
Wage sum, €						
2001	76,756	79,493	-0.007	-0.42	3,029	9,076
2000	69,862	76,164	-0.019	-1.29	2,894	8,620
1999	66,878	72,526	-0.018	-1.19	2,707	8,117
Wage sum growth						
2000–2001	8,652	4,500	0.018	0.96	2,761	8,231
1999–2000	4,096	5,857	-0.023	-1.38	2,597	7,810
1999–2001	13,459	9,824	0.014	0.74	2,540	7,664
Turnover						
2001	466,197	497,373	-0.009	-0.59	2,933	8,851
2000	464,644	474,579	-0.003	-0.16	2,894	8,620
1999	433,121	457,062	-0.007	-0.44	2,707	8,117
Turnover growth						
2000–2001	2,671	25,464	-0.024	-1.39	2,665	8,006
1999–2000	32,366	32,945	-0.001	-0.04	2,597	7,810
1999–2001	41,601	58,119	-0.018	-1.13	2,453	7,469
Operating profit						
2001	47,418	43,745	0.007	0.41	2,933	8,851
2000	57,472	42,665	0.015	0.83	2,894	8,620
1999	45,208	40,667	0.011	0.61	2,707	8,117
Operating profit growth						
2000–2001	-10,349	957	-0.021	-1.08	2,665	8,006
1999–2000	13,560	3,201	0.016	0.80	2,597	7,810
1999–2001	4,396	4,901	-0.003	-0.15	2,453	7,469

Notes: 1) Employment as in the Financial Statements data. 2) Measured in levels. All other growth variables are also in levels, not percentages. For the definition of the normalised difference and *t*-statistic, see Table 6.3.

The numbers in Table 6.4 are in levels but a reproduction of the table in logs yields qualitatively similar numbers, only with smaller (less significant) differences between the groups. A more rigorous way to look into the validity of the target-comparison grouping is to estimate the treatment effect model with dummy experiments for the pre-treatment years. This is done in the robustness checks section of the results chapter.

6.3 Data sets

The primary data sources are the company panel of Statistics Finland's Finnish Linked Employer Employee Data (FLEED), Business Register and the Structure of Earnings data. The FLEED company panel is compiled from the Financial Statements data and the information content is harmonised over the years. The company panel covers almost all active firms in Finland. The Business Register contains basic information on all establishments and firms.

The information on financial statements and balance sheets in the firm data come mainly from the tax authorities and are checked for consistency by Statistics Finland. The employment measure, the number of employees on a firm's payroll over the calendar year, which we used in our previous study, also came from the tax register. In this study, I use an alternative measure, the number of employees in the firm in the last week of the year. This is calculated from the FLEED employee panel⁸⁰ and was not previously available for the relevant years. I consider the cross section information on employment a more reliable measure of a firm's average annual employment than the tax register number. The Business Register data is used mainly to identify firms that reside entirely in either the target or comparison region of this study, i.e. that all establishments of a given firm are in the same area. That enables me to keep the multi-establishment firms in the data. There are only a few of those but as they are large firms, it is potentially important to keep them in the data instead of dropping them altogether. In principle, the firms with establishments both in the experiment and comparison regions would be very interesting cases but there are very few of those in the data and they are dropped from the sample.

The Structure of Earnings data come from Statistics Finland's data on wages and salaries, which is compiled by combining data collected by employer organisations from their members with those from Statistics Finland's own wage and salary survey. The Confederation of Finnish Industries (EK) collects comprehensive wage data from all of its member firms in October of each year. The data consists of complete payroll information, excluding top management and owners of the firms. EK member firms cover ~70% of Finnish GDP and have ~950,000 employees. The number of employees in EK member firms represents approximately half of entire private sector employment (~1.8 million in 2009). The Statistics Finland wage survey is sample-based and stratified by size category and industry classification. Wage and salary data on employees are collected from October. Only firms with five or more employees are sampled and

⁸⁰ The employee panel includes the total working age population in Finland. The firm panel also has an employment measure: average full-time equivalent yearly labour force. However, that number is imputed for most of the small firms and hence is not applicable in this study.

the sample covers ~10% of workers in unorganised firms. The Structure of Earnings wage data covers all organised employers and is *representative* of unorganised employers. However, while the firm-level data consists of near-complete firm populations, the Structure of Earnings data is much more limited in scope. The wage data covers ~5% of the target region firms and ~20% of the target region employees and the samples are by no means random. There is also quite a lot of yearly variation in the number of wage records per firm. Compared to the firms' personnel as calculated from the FLEED worker panel, or what is stated in the firm register, it seems that for some years much of the personnel of some firms is missing. Therefore, the validity of the wage data is not as good as the firm data. On the other hand, the quality of the information on wages should be very good, and much better than a proxy calculated from the firm data. As long as the method of selection into the dataset does not vary between the regions (it should not), comparisons should be possible. Information on hours is less accurately measured. It is calculated as 4.345 times the regular weekly hours plus overtime. The reporting of overtime varies, and for employees with a monthly salary (two thirds of the wage data) it might be a more error-prone measure than for the workers paid by the hour.

The datasets are available for research in the research laboratory of Statistics Finland.

6.4 Identification

The starting point for estimating the effect of the payroll tax reduction on firm- (or individual-) level responses y_i is a regression

$$y_{it} = c_i + \lambda_t + \tau w_{it} + \mathbf{x}_{it}\boldsymbol{\gamma} + u_{it}, \quad t = 1, \dots, T, \quad (21)$$

where λ_t are year effects, w_{it} indexes the treatment⁸¹, \mathbf{x}_{it} are the firm-level control variables, c_i is the firm fixed effect and u_{it} are the idiosyncratic errors⁸². Estimation by FE or first differencing to remove c_i is standard if the treatment is uncorrelated with u_{it} . Removing firm fixed effects would also remove any systematic differences between the treatment and comparison groups. While focusing on the changes differences away pre-existing dissimilarities between the target and comparison regions, it is still possible that the target and the comparison regions experience different shocks or display different pre-existing trends in the response variables. In particular, differing industrial structures may lead to different timing of the business cycle in the comparison and the target

⁸¹ w_{it} is a payroll tax cut indicator that is one if firm i gets the tax cut at time t and zero otherwise.

⁸² See Imbens and Wooldridge (2009), section 5, for a review of programme evaluation methods under unconfoundedness and section 6 for the selection in the unobservables case.

regions. It is easy to add region- or industry-specific time trends or their interactions to (21). It is also possible to account for differing trends for each firm:

$$y_{it} = c_i + g_i t + \lambda_t + \tau w_{it} + \mathbf{x}_{it} \boldsymbol{\gamma} + u_{it}, \quad t = 1, \dots, T. \quad (22)$$

Equation (22), a random linear trend model, is a special case of a correlated random coefficient model, which can be consistently estimated for $T \geq 3$ by first differencing

$$\Delta y_{it} = g_i + \eta_t + \tau \Delta w_{it} + \Delta \mathbf{x}_{it} \boldsymbol{\gamma} + \Delta u_{it}, \quad t = 2, \dots, T, \quad \text{where } \eta_t = \lambda_t - \lambda_{t-1} \quad (23)$$

and then running a fixed effects regression – or by differencing for a second time (Wooldridge, 2005).

If assignment to the treatment and comparison groups is a random draw or an unconfounded natural experiment, (23) estimated with standard regression methods will yield unbiased estimates and inference. Donald and Lang (2007), Bertrand, Duflo and Mullainathan (2004), and Hansen (2007a, 2007b) consider a case with unobserved group effects that introduce dependencies in the error terms between firms within groups (Donald and Lang) or over observations of the same units over time (Bertrand et al.), and how these could be dealt with in a setting where both the number of groups and observed time periods becomes large (Hansen). Not accounting for these group-wise or temporally correlated errors still gives the correct treatment effect estimate but invalidates inference.

In the Finnish tax cut case the number of groups is two or, at a stretch, three. Hence, the cluster sample methodology of Donald and Lang is not applicable. With two clusters, the cluster effect cannot be estimated and inference on the treatment effect estimator is impossible. I argue that in the Finnish case the test and comparison firms, although situated in geographically distinct areas, are actually in the same region as defined by firm subsidy rules. Therefore, it is unlikely that there would be a group effect large enough to swamp the sampling variance in sample means for treatment and comparison firms. Indeed, the identification (rather than the correct inference) of the tax cut effect hinges on the experiment being uncorrelated with other shocks in the target or comparison regions. On the other hand, it is likely that observations on the same firm are correlated over time. Therefore, I use one of the methods advocated by Bertrand *et al.* (2004) to take this type of error correlation into account.

In practice I first estimate (21) and (22) *without* the tax cut indicators. Then I aggregate⁸³ the error terms over the *target region* firms into pre- and post-treatment values and regress these on the treatment indicator. The treatment indicator is zero for 2001–2002 and one for 2003–2006 in Lapland and the indicator is zero over 2001–2004 and one over 2005–2006 in Kainuu. These are my preferred estimates reported in the next section.

I also ran regressions (21) and (22) directly, using policy change indicators for each year (2001,..., 2006) separately to better understand the timing of the effects and to see if experiments defined in this way obtain significant coefficients in wrong years. I comment on these and the other robustness checks in the next section. One direct observation is that estimating (21) is not sufficient: there appear to be trends in some of the response variables. Therefore, I report only results where trends are accounted for.

The original dataset has information on all relevant variables for the period 1999–2007. In the regression analysis, I use data on 2001–2006, i.e. from two years prior to the start of the experiment in Lapland until the experiment has run for two years in the Kainuu region. The main reason for doing this is to avoid using years too far from the tax change and thereby avoid mixing up the tax cut effect with other possible regionally occurring shocks. Another reason for dropping the year 2007 is the start of yet another regional employment subsidy scheme, where the experiment area partly overlaps with both the Kainuu region and the comparison region used in this study.

6.5 Results

The impact of the payroll tax cut is explored in this section. To account for the potentially heterogeneous effects of the tax cut I consider the results for four groups:

- 1) *all firms that existed*⁸⁴ *in 2001*, i.e. two years before the experiment started and before there was any common knowledge of the experiment,
- 2) *a group of firms where the most capital-intensive firms and the firms with the highest turnover per employee ratio (the firms in the highest quartile of either measure) are removed,*

⁸³ When the response variable is in levels, the aggregation is done by taking the mean of pre and post experiment residuals. In the case of first differenced responses, the aggregate is the sum of pre- and post-experiment residuals to capture the aggregate growth in the variables.

⁸⁴ I define existence as having positive turnover and wage sum.

- 3) firms where the part of payroll taxes to be deducted is well below ($\leq 25,000$) the maximum deduction limit, 30,000 euros a year, before the experiment starts (2001 and 2002) and hence face a lowered marginal labour cost and
- 4) the intersection of groups 2 and 3.

Group 1 is the base group and groups 2–4 are formed from it according to the above criteria.

These groupings are designed to focus on groups of firms intuitively the most sensitive to changes in labour costs and to ascertain that the possible effect of the tax cut is not drowned out by other strategic actions by large firms⁸⁵. Group 3 is probably the most interesting as in this group the tax cut makes hiring an extra employee cheaper and the restriction does not severely reduce the number of observations (see Table 6.5 for number of observations).

Table 6.5 Number of observations on the response variable Employment in firm groups 1–4 and number of firms.

Group	Number of observations				Number of firms			
	1	2	3	4	1	2	3	4
Levels								
Target	15,137	8,036	14,829	7,977	2,934	1,602	2,879	1,592
Control	45,483	26,576	44,451	26,262	8,807	5,257	8,617	5,200
Logs								
Target	12,518	6,605	12,214	6,549	2,706	1,456	2,651	1,446
Control	37,949	22,225	36,943	21,922	8,165	4,832	7,975	4,775
First differences and differences in relative changes[#]								
Target	14,953	7,922	14,646	7,863	2,934	1,602	2,879	1,592
Control	44,904	26,197	43,877	25,885	8,801	5,252	8,611	5,195
Differences in logs								
Target	11,750	6,226	11,447	6,170	2,566	1,375	2,511	1,365
Control	35,710	20,972	34,710	20,672	7,787	4,603	7,598	4,546

[#]) Relative changes calculated as $P = \frac{X_t - X_{t-1}}{\frac{1}{2}(X_t + X_{t-1})}$, i.e. $P \in [-2, 2]$. $P = 0$ when X is zero for periods $t - 1$ and t .

⁸⁵For example, one large electronic components supplier in Lapland shifted its entire operation to China, resulting in a large employment effect, certainly not related to the experiment.

Some of the response variables are not defined for all observations: there are no logarithms for non-positive values, relative changes⁸⁶ are not defined for two consecutive missing values (for two consecutive zeros I set the relative change to zero) and for differences one needs two consecutive non-missing observations. I choose not to limit the observations to those where all the responses exist and therefore, in addition to the groupings, the number of observations differs across the response variables. This decision does not affect the estimates much but helps to tighten the confidence intervals by making use of all available information. I report the number of observations on each response type for employment for the aforementioned groupings in Table 6.5. The numbers for the other response variables follow these closely and are not reported.

The most important difference between the measures (levels, logs, differences, differences in logs, differences in relative changes) is that for the levels, differences and differences in relative changes I have added an observation for firm exits. Otherwise, the last change in e.g. employment from a positive value to zero would be omitted. For logs or changes in logs, this is not possible.

6.5.1 Where did the money go?

As I stated already in the introduction, I found no clear effects on any of the measures of firms' performance. The question arises then whether the experiment is too small to register in the data at all or if the quality of the data is too poor for the purpose. In this subsection I go directly after the tax cut.

Compliance on the part of the firms is not a problem. Taking part in the experiment only requires firms to notify their local tax office that the firm is not going to pay the first 30,000 euros of the combined national pension insurance and national health insurance and to report the deducted amount at the end of the calendar year. According to the data from the tax authorities, practically all the firms with employees in the target area filed a starting declaration.

The most disaggregated measure in the data containing the waived part of payroll taxes (national pension insurance and national health insurance) is called "other labour costs". These labour costs are directly related to the wage sum, excluding pension contributions. The waived part is a little over half of this "other labour costs" entry. The other half consists mainly of accident, unemployment and group life insurance payments (see Table 6.1). The exact amount of the reduced payroll taxes for each firm would be available from the tax records, but in order to see how reliable the firm register data is, I estimated the effect of the payroll tax reduction on other labour costs. If it were not visible at all, it would raise a

⁸⁶ I use the definition introduced e.g. in Davis, Haltiwanger and Schuh, 1996, see notes under Table 6.5.

serious concern about data quality. The results of these regressions are presented in Table 6.6. The first four *columns* contain estimates from regressions without firm-specific slopes, but there are controls for region and industry trends. Columns five to eight are from regressions with firm-specific slopes.

Table 6.6 *Effect of tax cut on firms' "other labour costs"*

Group	1	2	3	4	Firm fixed effects			
					1	2	3	4
First differences	-527 (341)	-588 (187)	-453 (159)	-480 (168)	-711 (307)	-619 (176)	-525 (146)	-490 (159)
Differences in logs	-0.1519 (0.0363)	-0.0996 (0.0521)	-0.1506 (0.0370)	-0.1002 (0.0524)	-0.1503 (0.0327)	-0.0957 (0.0457)	-0.1476 (0.0332)	-0.0931 (0.0460)
Differences in relative changes [#]	-0.1477 (0.0333)	-0.1560 (0.0485)	-0.1456 (0.0339)	-0.1547 (0.0488)	-0.1611 (0.0308)	-0.1719 (0.0445)	-0.1576 (0.0313)	-0.1678 (0.0446)
Differences in share of wage sum	-0.0232 (0.0035)	-0.0006 (0.0045)	-0.0232 (0.0035)	-0.0004 (0.0045)	-0.0227 (0.0029)	-0.0002 (0.0043)	-0.0228 (0.0030)	0.0000 (0.0043)

Notes: Coefficients in **bold** are significant at the 5% risk level, standard errors in (parenthesis). Columns marked as follows: 1) all firms functional in 2001, 2) firms where turnover / employee ratio and capital intensity are in the highest third are dropped, 3) firms with potential payroll tax cut $\leq 25,000$ € up till 2002, 4) firms fulfilling conditions 3 and 4.

#) Relative changes are calculated as $P = \frac{X_t - X_{t-1}}{\frac{1}{2}(X_t + X_{t-1})}$, i.e. $P \in [-2, 2]$. $P = 0$ when X is zero for periods

$t - 1, t$. All regressions have controls for year effects and firm age. In all regressions without firm fixed effects industry is controlled at the 4-digit level (338 classes in data) to account for industry-specific trends.

The "other labour cost" regressions show clear reductions, but the reductions are smaller than what one would expect. The diff-in-diff's estimates are approximately 600 euro, or 2.3% of the wage sum. What is notable, however, is that in groups 2 (and 4) the effects in terms of share of the wage sum are very small or non-existent. A more careful inspection of the one-year treatment effect regressions shows that there might be a problem with the comparability of the treatment and comparison regions after all: many response variables for the target region firms for 2006, especially in Lapland, indicate that some other factors besides the tax experiment are probably driving the differences between the regions. Hence I re-estimated the models for differences so that the treatment in Lapland lasted only for two years (as in Kainuu), instead of four. The results of this exercise are given in Table 6.7. Now almost all the effects on other labour costs become larger and more statistically significant – in particular the results for groups 2 and 4 change.

Table 6.7 Tax cut effect on firms' "other labour costs", 2-period treatment

Group	1	2	3	4	Firm fixed effects			
					1	2	3	4
First differences	-1,264 (343)	-870 (195)	-1,033 (172)	-783 (176)	-1,311 (297)	-867 (180)	-1,044 (160)	-783 (163)
Differences in logs	-0.3783 (0.0393)	-0.2568 (0.0543)	-0.3770 (0.0401)	-0.2543 (0.0547)	-0.3573 (0.0369)	-0.2172 (0.0498)	-0.3554 (0.0376)	-0.2133 (0.0501)
Differences in relative changes [#]	-0.3340 (0.0353)	-0.2826 (0.0497)	-0.3291 (0.0359)	-0.2790 (0.0500)	-0.3354 (0.0330)	-0.2833 (0.0462)	-0.3301 (0.0335)	-0.2779 (0.0464)
Differences in share of wage sum	-0.0205 (0.0033)	-0.0118 (0.0045)	-0.0202 (0.0034)	-0.0117 (0.0045)	-0.0193 (0.0030)	-0.0112 (0.0042)	-0.0192 (0.0030)	-0.0112 (0.0043)

Notes: see Table 6.6 for explanatory notes.

However, dropping treatment status from the years 2005 and 2006 in Lapland does not change the results for any other responses – none of the firm-level responses becomes statistically significant and the possible negative effect on wages remains the same.

6.5.2 Effect on employment, wage sum and profits

The coefficients of the treatment indicator from regressions where the dependent variable is employment are reported in Table 6.8. The estimates show no statistically significant effects on employment. All estimates where the unit of measurement is employees (row 1) are negative, meaning that the aggregate effect for the target region was also negative. The estimates are positive when the response is measured in differences-in-differences in log employment (without firm-specific trends) and diff-in-diff's in percentage terms (with and without firm-specific trends). The differences in growth estimators in the lowest row are actually not that small and show the largest "effect" for group 3, but the standard errors are far too large to warrant any conclusions regarding positive effects. While the differences in logs and differences in relative changes measure the same thing, the results differ owing to the exclusion of zero employment observations from the logs.

Table 6.8 *Effect of tax cut on employment in firms*

Group	1	2	3	4	Firm fixed effects			
					1	2	3	4
First differences	-0.019 (0.107)	-0.039 (0.105)	0.034 (0.078)	-0.071 (0.101)	-0.077 (0.096)	-0.049 (0.101)	-0.045 (0.075)	-0.081 (0.098)
Differences in logs	0.0017 (0.0178)	0.0149 (0.0224)	0.0058 (0.0178)	0.0140 (0.0226)	-0.0141 (0.0166)	-0.0017 (0.0209)	-0.0102 (0.0166)	0.0000 (0.0208)
Differences in relative changes [#]	0.0338 (0.0297)	0.0142 (0.0386)	0.0369 (0.0302)	0.0122 (0.0388)	0.0127 (0.0273)	0.0066 (0.0353)	0.0152 (0.0277)	0.0070 (0.0354)

Notes: Coefficients in **bold** are significant at 5% risk level, standard errors in (parenthesis). Columns marked as follows: 1) all firms functional in 2001, 2) firms where turnover / employee ratio and capital intensity are in the highest third are dropped, 3) firms with potential payroll tax cut $\leq 25,000$ € up till 2002, 4) firms fulfilling conditions 3 and 4.

#) Relative changes are calculated as $P = \frac{X_t - X_{t-1}}{\frac{1}{2}(X_t + X_{t-1})}$, i.e. $P \in [-2, 2]$. $P = 0$ when X is zero for periods

$t-1, t$. All regressions have controls for year effects and firm age. In all regressions without firm fixed effects, industry is controlled at the 4-digit level (338 classes in data) to account for industry-specific trends.

Findings from our previous study⁸⁷ and the findings from other recent studies of northern Sweden and Norway showed that a payroll tax cut is likely to push wages up. In Table 6.9 I report the results from regressions on firms' yearly wage sum.

Table 6.9 *Effect of tax cut on firms' wage sum*

Group	1	2	3	4	Firm fixed effects			
					1	2	3	4
First differences	2,732 (2,268)	56 (1,660)	2,427 (1,272)	-345 (1,508)	1,720 (2,038)	-29 (1,527)	1,425 (1,122)	-414 (1,309)
Differences in logs	0.0134 (0.0144)	0.0140 (0.0175)	0.0117 (0.0145)	0.0117 (0.0175)	0.0014 (0.0135)	-0.0007 (0.0165)	0.0009 (0.0136)	0.0001 (0.0166)
Differences in relative changes [#]	0.0306 (0.0202)	0.0087 (0.0284)	0.0300 (0.0205)	0.0086 (0.0286)	-0.0007 (0.0173)	-0.0214 (0.0245)	-0.0002 (0.0175)	-0.0192 (0.0244)

See Table 6.8 for explanatory notes.

⁸⁷ We did not find any statistically significant effects on wage sum but some indication that wage rates had risen in service industries. Our earlier estimate for the wage effect, 1,728 euros, was a diff-in-diff's five nearest neighbours matching estimator for Lapland for the years 2003 and 2004. Curiously enough, here the diff-in-diff's estimator with firm-specific slopes for Lapland (2003–2006) and Kainuu (2005–2006) is a very close hit: 1,720 euros.

None of the estimates in Table 6.9 is statistically significantly non-zero. The aggregate effect on the target area wage sum is positive (group 1, row 1). The estimate coming closest to being significant is the diff-in-diff's estimator (2,427 euro) for group three without firm-specific trends. The differences in relative changes estimators that were positive for employment are also positive here, but are smaller.

Compared to the wage and employment effects, the detection of profit effects is made even harder by the fact that profits is a quantity containing far more idiosyncratic and time series variation than the wage sum or employment. The measure for profits I use is operating profit. Due to changes in accounting practises, this is the only profit measure in the data that is consistent over time. As operating profits quite often show negative values (23% of the observations), the taking of logs and calculating relative changes is not very meaningful. Therefore, I took proportional measures of profits relative to the wage sum. Operating profit relative to the wage sum does not directly measure the effect of the tax cut on profits but gives an indication of whether the tax cut had an effect over and above the effect on wages. The estimation results in Table 6.10 show, again, no statistically significant coefficients.

Table 6.10 Effect of the tax cut on firms' operating profit

Group	1	2	3	4	Firm fixed effects			
					1	2	3	4
First differences	8,951 (19,823)	1,703 (1,921)	1,283 (1,792)	-297 (1,436)	9,811 (20,353)	2,077 (1,895)	395 (1,856)	27 (1,377)
Difference in the share of wage sum	0.0265 (0.0297)	-0.0140 (0.0276)	0.0265 (0.0298)	-0.0147 (0.0278)	0.0223 (0.0277)	0.0053 (0.0286)	0.0232 (0.0279)	0.0048 (0.0288)

See Table 6.8 for explanatory notes.

6.5.3 Effect on wages and hours, individual wage records

The number of individual wage and hours observations is decent (see Table 6.11) but the number of firms is small compared to the total number of firms. I have maintained the same grouping (1–4) as in the previous subsection. The added groups are

5) observations where a worker stayed in the same firm and occupation from $t-1$ to t and

6) firms on 4-digit industry common support, i.e. in industries that are found in both the treatment and comparison areas.

The main reason for forming group 5 is to reduce noise. On the other hand, it could be argued that wage changes often occur in conjunction with a change of occupation or employer and therefore such movers should be kept in the data. This grouping might also be used to account for different mobility patterns across the regions but such differences do not exist. I generated group 6 to account for the fact that in the wage data there are a few firms operating in industries that were completely lacking in the target (or control) area and thus one could consider these parts of the data incomparable to the other region, even after controlling for industry.

Table 6.11 Number of observations on the response variable hourly wages in worker-firm groups 1–6. Also a number of individual workers and distinct firms in the wage data.

Group	1	2	3	4	5	6
N observations						
Levels & logs						
Target	9,556	2,110	4,271	1,256	4,473	8,270
Control	38,016	9,916	14,379	5,539	19,910	23,547
Diff's, diff's in logs, Differences in relative changes						
Target	4,951	988	2,118	598	4,095	4,142
Control	21,356	4,796	7,157	2,444	18,746	12,897
N individuals						
Levels & logs						
Target	4,004	965	1,973	632	1,944	3,595
Control	14,846	4,636	6,584	2,874	8,022	9,411
Diff's, diff's in logs, Differences in relative changes						
Target	2,013	451	946	287	1,813	1,731
Control	8,242	2,058	3,069	1,110	7,675	5,131
N firms						
Levels & logs						
Target	179	49	151	46	132	158
Control	561	211	450	180	424	412
Diff's, diff's in logs, Differences in relative changes						
Target	135	38	109	35	127	122
Control	412	149	322	123	405	311

The evidence from the wage regressions to back up the earlier result of the tax cut being channelled to higher wages is scant. I present the results in Table 6.12. The differences-in-differences estimates without worker fixed effects (left half of Table 6.12) are mostly negative and also statistically significant for group 6, workers in the industry common support, indicating a wage drop of 26 cents per

hour, or 1.6 percent. When worker-specific slopes are added, all estimates become smaller in absolute value but those for group 6 retain their statistical significance. Here the standard errors are also tight enough to show that wage changes cannot have accommodated any large changes in the firms' wage sum.

Table 6.12 Effect of tax cut on hourly wages

Group	Worker fixed effects											
	1	2	3	4	5	6	1	2	3	4	5	6
First differences												
	-0.14	-0.03	-0.03	0.35	-0.19	-0.26	-0.11	-0.23	-0.05	-0.15	-0.10	-0.21
	(0.09)	(0.35)	(0.17)	(0.59)	(0.09)	(0.11)	(0.08)	(0.19)	(0.10)	(0.30)	(0.07)	(0.09)
Differences in logs												
	-0.0049	-0.0146	-0.0125	0.0010	-0.0072	-0.0162	-0.0020	-0.0089	-0.0002	0.0005	-0.0021	-0.0117
	(0.0059)	(0.0159)	(0.0090)	(0.0205)	(0.0060)	(0.0069)	(0.0046)	(0.0101)	(0.0060)	(0.0147)	(0.0043)	(0.0053)
Differences in relative changes												
	-0.0053	-0.0146	-0.0124	0.0003	-0.0074	-0.0159	-0.0020	-0.0085	0.0001	0.0012	-0.0020	-0.0115
	(0.0058)	(0.0154)	(0.0087)	(0.0196)	(0.0058)	(0.0067)	(0.0045)	(0.0099)	(0.0058)	(0.0143)	(0.0042)	(0.0052)

Notes: Coefficients in **bold** are significant at the 5% risk level, standard errors in (parenthesis). Columns marked as follows: 1) all firms functional in 2001 (positive employment and wage sum), 2) firms where turnover / employee ratio and capital intensity are in the highest third are dropped, 3) firms with potential payroll tax cut $\leq 25,000$ € up till 2002, 4) firms fulfilling conditions 2 and 3, 5) observations where worker stayed in the same firm and occupation from $t-1$ to t , 6) firms on 4-digit industry common support, i.e. in industries that are found in the treatment and comparison areas.

All regressions have controls for year effects. Industry is controlled up to 16 classes. The individual controls are gender, education level, age and age squared, tenure and indicators for firm or occupation change. Occupation is controlled at the 3-digit level (91 classes in the data).

The wage sum could rise more than employment without changes in hourly wages if the hours worked increase. The Structure of earnings data has information on monthly hours but how accurately it measures the actual hours worked varies across industries and depends on which collective agreement is followed. The regression results where the dependent variable is “monthly hours worked” are presented in Table 6.13. The only statistically significant estimates are found for group 4 when the measure is diff-in-diff's without employee-specific slopes. Most of the estimates measuring proportional changes (differences in logs, differences in relative changes) are negative.

Table 6.13 Effect of tax cut on monthly hours*

Group	Worker fixed effects											
	1	2	3	4	5	6	1	2	3	4	5	6
First differences	0.51	2.74	1.07	5.40	0.50	0.02	0.28	0.59	0.63	2.29	0.34	-0.13
	(0.72)	(1.67)	(1.23)	(2.30)	(0.69)	(0.82)	(0.50)	(0.97)	(0.79)	(1.56)	(0.45)	(0.57)
Differences in logs	-0.0002	-0.0050	-0.0047	-0.0023	-0.0011	-0.0026	-0.0002	-0.0012	0.0002	0.0002	-0.0007	-0.0009
	(0.0019)	(0.0061)	(0.0026)	(0.0075)	(0.0019)	(0.0021)	(0.0012)	(0.0034)	(0.0016)	(0.0048)	(0.0011)	(0.0014)
Differences in relative changes [#]	-0.0003	-0.0049	-0.0046	-0.0024	-0.0011	-0.0026	-0.0002	-0.0010	0.0002	0.0004	-0.0007	-0.0008
	(0.0019)	(0.0059)	(0.0026)	(0.0072)	(0.0019)	(0.0021)	(0.0012)	(0.0033)	(0.0015)	(0.0047)	(0.0011)	(0.0013)

Notes: Coefficients in **bold** are significant at the 5% risk level, standard errors in (parenthesis). See Table 6.12 for description of the grouping.

* Hours are calculated as 4.345 times the regular weekly hours + overtime.

6.5.4 Robustness checks

The payroll tax experiment did not, on average, have a *statistically significant* effect on employment, firms' wage sum or operating profit at a regional level. It might have had a surprisingly negative effect on hourly wages. I did some robustness checks to scrutinise these results.

The role of firm exits

The proportional diff-in-diff's estimator for employment is larger than the changes in log's estimator (Table 6.8). As the changes in log's estimator omits the effect of firm exits, the result gives an indirect indication that the tax break might have helped some firms to continue operating rather than exiting. The yearly number and share of exits of firms that existed in 2001 are given in Table 6.14. The definition of exit is based on information in Statistics Finland's firm and establishment registers: a real exit occurs during a year if the firm code and all establishments linked to it at the end of the previous year disappear from the register.

Table 6.14 Share and number of firm exits* in target and comparison regions

	Lapland		Kainuu		Comparison region	
	Share of exits	# of exits	Share of exits	# of exits	Share of exits	# of exits
2002	0.0267	39	0.0357	51	0.0301	261
2003	0.0273	38	0.0407	55	0.0249	204
2004	0.0248	33	0.0317	40	0.0238	186
2005	0.0252	32	0.0319	38	0.0257	191
2006	0.0284	34	0.0264	30	0.0274	194

* Exits are defined based on firm and establishment registers. An exit has occurred if the firm identifier and all establishments linked to it disappear from the register in the year following the potential exit. The shaded area in the table indicates the tax experiment.

I estimated the effect of the tax experiment on exits using a Cox proportional hazards model where I controlled for the treatment area (Lapland and Kainuu separately), industry and the firm's age group. The coefficient of interest is the tax cut indicator. I used first the same dataset as in the regressions on employment. Then I estimated the model for an extended time period where I included all firms functional in 1999 in order to have more pre-experiment years for Lapland. The effects of the tax cut on firm exits are similar and they are not statistically significant in either case. I report the results only for the 2001–2006 period (see Table 6.15).

Table 6.15 Effect of tax cut on firms' exit probability. Odds ratios from proportional hazards model.

Group	1	2	3	4
	0.9193	1.0382	0.9090	1.0253
	(0.1097)	(0.1486)	(0.1093)	(0.1469)

Notes: Odds ratios in bold are significant at the 5% risk level, standard errors in (parenthesis). Columns marked as follows: 1) all firms functional in 2001, 2) firms where turnover / employee ratio and capital intensity are in the highest third are dropped, 3) firms with potential payroll tax cut $\leq 25,000$ € up till 2002, 4) firms fulfilling conditions 3 and 4.

The use of a stock sample, i.e. firms that exist in a certain time period, would be problematic if my interest was in the duration dependence of firm survival. Here the focus is on the effect of the tax cut on the existing firm population and hence the oversampling of long-standing firms would be a different issue.

Other subgroups

The results were derived for subgroups that I considered sensible but that were somewhat arbitrary. Here I comment on the results for a number of groupings, omitting the tables for the sake of brevity. The first group comprises firms that existed throughout the entire observation period of 2001–2006. This is a potentially more stable group of firms, and a group that was exposed to the experiment for the longest time span. This kind of constraint also excludes exiting firms from the sample. On the one hand, exits most definitely belong to the data as a vital part of firm dynamics. On the other hand, exits of large firms could have a substantial effect on total employment in the region and it could be argued that relocating a manufacturing firm to China or Estonia is not directly related to small changes in payroll taxes. Therefore excluding these events would be justified. These results, however, are very close to those obtained for group 1. If one were to look very carefully, the drop in the other labour cost variable is largest and the rise in the wage sum biggest in this group, but the differences are small.

To account for the fact that the wage regressions were run on a small subgroup of the firm data, I also ran the firm-level regressions for the same subgroup of firm-year observations that occur in the wage data. The possibly negative wage effect is indeed mirrored in these results. Most of the coefficients in the wage sum regressions are negative but do not differ significantly from zero. The estimates of changes in operating profit are much larger than for the complete firm data. Owing to the small sample size, none of the effects is statistically significant.

For two thirds of the workers in the wage data, hourly pay is calculated from monthly wages and information on actual hours worked. If the information on hours were for some reason less reliably recorded than wages for this group, it could lead to a blurring of the results. I ran separate regressions for workers with hourly wage and monthly salary and the results do not reveal any large differences. The wage effects (still mostly insignificant) are, however, consistently more negative for hourly paid workers and the monthly hours effects are larger for salary earners. The difference between the coefficients is insignificant.

The small firms (group 3) were defined in terms of their pre-experiment payroll tax payments. Another way of defining small firms is to use a clear personnel limit. If I apply a limit of a maximum of 10 employees on any of the pre-experiment years, I capture 90% of the firms. They account for almost exactly half of the total employment in the target and comparison area firms. The results for this group are very close to the group 3 results.

The last subgroup I considered was the large firms, defined as being above the maximum payroll tax cut for both 2001 and 2002. This is a small group and does

not weigh too much in the firm-level regressions. None of the treatment effects is statistically significant for them alone. One reassuring finding amidst the non-existent effects was that the differences-in-differences estimate with firm-specific slopes for the reduction in other labour costs is 23,483 euros for the basic pre-post estimate and 28,539 euros when the treatment is limited to the two post-experiment years in Lapland. This gives an indication that the data correctly captures the maximum tax cut for this group of firms.

Yearly (placebo) experiments

In my regression set-up, the last stage regression is run on observations of one pre- and one post-experiment observation per firm. As there are several pre- and post-experiment years in the original data, it makes it possible to estimate a yearly effect separately for each year. If statistically significant differences appear even before the experiment started, this could indicate that something other than the payroll tax experiment is driving the results.

Employment. None of the yearly effects is statistically significant. There also does not seem to be any pattern to these effects, neither in the aggregates nor if I look at Lapland and Kainuu separately.

Wage sum. A negative effect on the *level* of the wage sum can be seen to originate from a downward trend in treatment area wage sums compared to the comparison area. The estimates for the first observation years are positive and the estimates decline year by year. None of the one-year estimates for first differenced responses is statistically significant.

Operating profit. None of the one-year effects is statistically significant. The diff-in-diff's-type estimates are greatest and positive for the start year of the experiment (2003 in Lapland and 2005 in Kainuu), giving some weak evidence that the firms might have pocketed the savings.

Hourly wage. The significant and negative effect obtained for the industry common support group (group 6 in Table 6.12) originates from the year the experiment started in Kainuu (2005), where there was a clear overall wage drop of three and a half per cent. In Lapland, however, there is a positive effect of 1.7 per cent in 2003 – after a negative and significant effect of 3.4 per cent in 2002. The other one-year effects are not statistically significant. That the negative effect occurs only in Kainuu and that there are statistically significant effects in “wrong” years implies that the observed negative effect has probably more to do with Kainuu-specific firm dynamics than with the payroll tax experiment.

The overall conclusion from the yearly effect exercise is that a) in most cases the standard errors are large compared to the pre-post regressions and b) the results for profits and wage sum give (very weak) support to the previous findings in the sense that the timing of the effects seems correct and c) the negative trend in the

level of the wage sum implies that the firm- or region-specific trends are important and the fixed effects regressions on levels and logs are mis-specified.

Heterogeneous treatment

The size of the tax cut depends on the payroll tax class of the firm and could vary from three to six per cent of the wage sum, up to the maximum reduction of 30,000 euros. It would therefore appear reasonable to use the actual reduction percentage as the treatment, not just a dummy for being in the target region. Almost all of the small firms (groups 2–4) belong to the lowest payroll tax category, however, and even though I can classify the larger firms quite accurately into the right categories, there are bound to be some errors. That makes the usefulness of differentiated treatment level assignment less useful than it first appears. The actual percentage reduction in payroll taxes is also endogenous for the large firms, where the tax cut is topped at 30,000 euros, since it depends on the observed wage sum. I ran the regressions with the percentage reduction as the treatment but the results were almost identical with those obtained with a treatment indicator.

Lapland and Kainuu as separate experiments

I estimated models separately for the Lapland and Kainuu regions, keeping the comparison region constant. Keeping in mind the results for the “other labour costs” regression indicating that the year 2006 might be a problem in Lapland, I also estimated two-year treatment effects for Lapland.

The differences between the results concerning employment are small. The estimates for employment effects are close to each other for Lapland and Kainuu and none is statistically significant. The effect on the wage sum is larger in Kainuu but again not statistically significant, whereas the estimates for changes in operating profit are larger in Lapland, still without being significantly non-zero. The hourly wage regressions reflect the results already reported in the one-year effects section; there is a negative and significant effect in Kainuu and a positive and significant effect in Lapland that cancel each other out, yielding no overall effect. This is the largest difference in results compared to the Korkeamäki and Uusitalo paper.

Seasonality of employment and the data

The employment measure for a firm used in this study is the number of workers with an employment relationship with it at the end of the year. The wage sum and profit measures are for the accounting period, which is most often the calendar year. The wage records are for the October of each year. This time pattern could hide some effects in wages and employment. Spring and summer are the high seasons for tourism in Kainuu and Lapland and one could argue that a temporary reduction in labour costs would have the largest effect on short-term

work contracts during the high season. Unfortunately, municipality-level monthly employment figures are not available. The only source would be the Labour force survey, but there the sample size in the treated region is not large enough for this purpose. Quarterly employment figures are available to a decent approximation of the target and comparison areas, however, and they show that employment is at its lowest in the first quarter, then rises 5–9 per cent for the middle quarters, after which it drops in the last quarter to close to the first quarter value. The pattern across regions is stable and similar and does not give rise to any concerns that seasonality masks some effects from the experiment.

Selecting the observation period

Dropping years from the beginning or the end of the current observation period of 2001–2006 has little effect on the results. The aforementioned trouble with 2006 in Lapland is to some extent present for Kainuu as well. Dropping that year lowers all the estimates a little, although none change their sign. Dropping the year 2001 has the opposite effect, i.e. the effects become somewhat larger but all remain well under the limit of becoming statistically significant.

6.6 Discussion and some conclusions

The main results of this study are that the payroll tax cut did not have a statistically significant effect on total employment, the wage sum, profits or wages in the target area. Most of the estimates are positive but unfortunately the standard errors are so wide that they could accommodate values indicating full shifting of the tax cut to either the wage sum, profits or, indeed, to employment. If we look at the euro-valued point estimates, the conclusion would be that the wage sum in the target region firms rose, employment did not and profits grew the most. Alternatively, if we consider the point estimates of the percentage changes, employment and the wage sum did grow by an equal amount and profits did not react. The only unambiguous finding is that the tax cut can be found from the Financial Statements data, although even there there was some uncertainty in the case of small and the least capital-intensive firms.

The effect on hourly wages found in Korkeamäki and Uusitalo (2009) is not found here for the combined target region of Kainuu and Lapland. The effect is still found for Lapland – but the estimates for Kainuu would imply a negative wage effect. The results also show one statistically significant change in a non-experiment year and hence do not warrant any strong conclusions.

Irrespective of the findings from this and other similar experiments in the Nordic countries, national pension insurance payments have been gradually lowered over recent years. From the beginning of 2010 they were abolished altogether, on the grounds that it would be beneficial for employment. There was some debate if

this was the most effective way to help firms generate jobs, but empirical facts had a rather small role in the discussion. This is partly due to a lack of such facts.

The Finnish payroll tax experiment is a rare example where a tax change is made in an experimental setting with the stated aim of facilitating economic research. Hence it is important to evaluate it, even if the results tell rather little about the effects on employment. This is also an opportunity to gather information on the experiment itself to learn more about how possible future experiments should be designed and implemented to the greatest scientific advantage. Based on my results I argue that it is still important to continue experimenting – it is also important to pre-evaluate future experiments to see if they are likely to yield accurate and reliable results.

References

- Aaronson, D. (2001): Price pass through and the minimum wage. *Review of Economics and Statistics*, 83, pp. 158–169.
- Aaronson, D. – French, E. (2007): Product market evidence on the employment effects of the minimum wage. *Journal of Labour Economics*, 25(1) pp. 167–200.
- Anderson, P. M. – Meyer, B. D. (2000): The effects of the unemployment insurance payroll tax on wages, employment, claims and denials. *Journal of Public Economics* 78(1–2), pp. 81–106.
- Angrist, J. D. – Pischke J. (2009): Mostly Harmless Econometrics – An Empiricist’s Companion, Princeton University Press.
- Benmarker, H. – Mellander, E. – Öckert, B. (2009): Do regional payroll tax reductions boost employment? *Labour Economics* 16, pp. 480–489.
- Bertrand, M. – Duflo, E. – Mullainathan, S. (2004): How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics* 119, pp. 249–275.
- Borjas, G. (2005): Labor Economics. McGraw-Hill (3rd edition).
- Brown, C. (1999): Minimum Wages, Employment, and the Distribution of Income. Chapter 32 in Ashenfelter, O. – Card, D. (eds.): *Handbook of Labour Economics*, North Holland Press.
- Card, D. – Krueger, A. B. (1995): Myth and measurement: the new economics of the minimum wage. Princeton University Press.
- Davis, S. J. – Haltiwanger, J. – Schuh, S. (1996): Job Creation and Destruction. The MIT Press, Cambridge.

- Donald, S. G. – Lang, K. (2007): Inference with difference-in-differences and other panel data. *The Review of Economics and Statistics*, 89(2), pp. 221–233.
- Draca, M. – Machin, S. – Van Reenen, J. (2008): Minimum wages and firm profitability. NBER Working Paper 13996, National Bureau of Economic Research.
- Flinn, C. J. (2006): Minimum wage effects on labor market outcomes under search, matching, and endogenous contact rates. *Econometrica*, 74(4), pp. 1013–1062.
- Goerke, L. (1996): Taxes on payroll, revenues and profits in three models of collective bargaining. *Scottish Journal of Political Economy*, 43(5), pp. 549–565, November.
- Hansen, C. B. (2007a): Asymptotic properties of a robust variance matrix estimator for panel data when T is large. *Journal of Econometrics* 141, pp. 597–620.
- Hansen, C. B. (2007b): Generalized least squares inference in panel and multilevel models with serial correlation and fixed effects. *Journal of Econometrics* 140, pp. 670–694.
- Imbens, G. W. – Rubin, D. B. (forthcoming): Causal inference in statistics and the social sciences. Cambridge University Press, Cambridge and New York.
- Imbens, G. W. – Wooldridge, J. M. (2009): Recent developments in the econometrics of program evaluation. *Journal of Economic Literature*, 47(1), pp. 5–86.
- Korkeamäki, O. – Uusitalo, R. (2009): Employment and wage effects of payroll tax cut – evidence from a regional experiment. *International Tax and Public Finance*, 16 pp. 753–772.
- Korkeamäki, O. (2009): The Finnish payroll tax cut experiment revisited. VATT Working Papers 22, Government Institute for Economic Research, Helsinki.
- Laki Kainuun hallintokokeilusta 9.5.2003/343 (Law of the Kainuu self-government experiment).
- Machin, S. – Manning, A. – Rahman, L. (2003): Where the minimum wage bites hard: the introduction of the UK National Minimum Wage to a low wage sector. *Journal of the European Economic Association*, 1, pp. 154–180.
- Murphy, K. J. (2007): The impact of unemployment insurance taxes on wages. *Labour Economics* Vol. 14(3), pp. 457–484.
- Neumark, D. – Wascher, W. (2007): Minimum wages and employment. *Foundations and Trends in Microeconomics*, 3(1–2), pp. 1–182.

Stewart, M. (2004): The impact of the introduction of the UK minimum wage on the employment probabilities of low wage workers. *Journal of the European Economic Association*, 2, pp. 67–97.

Wooldridge, J. M. (2005): Fixed effects and related estimators for correlated random-coefficient and treatment effect panel data models. *Review of Economics and Statistics* 87(2), pp. 385–390.

VATT JULKAISUT / VATT PUBLICATIONS

VATT JULKAISUT SARJASSA AIEMMIN ILMESTYNEET JULKAISUT

41. Verotus, talouspolitiikka ja kansantalous – Taxation, Economic Policy and the Economy. Vesa Kannianen – Seppo Kari (toim.). Helsinki 2005.
42. Suomi 10 vuotta Euroopan unionissa – Taloudelliset vaikutukset. Jaakko Kiander – Antti Romppanen (toim.). Helsinki 2005.
43. Suomi 2025 – Kestävän kasvun haasteet. Reino Hjerppe – Juha Honkatukia (toim.). Helsinki 2005.
44. Kasvumallin muutos ja veropolitiikan haasteet. Reino Hjerppe – Seppo Kari – Jaakko Kiander (toim.). Helsinki 2006.
45. Finance and Incentives of the Health Care System. Proceedings of the 50th Anniversary Symposium of the Yrjö Jahnsson Foundation. Antti Suvanto – Hannu Vartiainen (Eds.). Helsinki 2007.
46. Hyvinvointipalvelujen tuottavuus: Tuloksia opintien varrelta. Aki Kangasharju (toim.). Helsinki 2007.
47. Vaikuttavaa tutkimusta – miten arviointitutkimus palvelee päätöksenteon tarpeita? Seija Ilmakunnas – Teuvo Junka – Roope Uusitalo (toim.). Helsinki 2008.
48. Hyvinvointipalveluja entistä tehokkaammin. Uudistusten mahdollisuuksia ja keinoja. Seija Ilmakunnas (toim.). Helsinki 2008.
49. Terveyskeskusten tehokkuutta ja tuottavuutta selittävät tekijät. Juho Aaltonen – Maija-Liisa Järviö – Kalevi Luoma (toim.). Helsinki 2009.
50. Essays on globalization – Policies in trade, development, resources and climate change. Leena Kerkelä. Helsinki 2009.
51. Essays on Migration. Matti Sarvimäki. Helsinki 2009.
52. Essays on income inequality, poverty and the evolution of top income shares. Marja Riihelä. Helsinki 2009.
53. Essays on the efficiency of schools and student achievement. Tanja Kirjavainen. Helsinki 2009.
54. Verotuksen ja sosiaaliturvan uudistaminen – miksi ja mihin suuntaan? Essi Eerola – Seppo Kari – Jaakko Pehkonen (toim.). Helsinki 2009.
55. Talouden rakenteet 2009. Seija Ilmakunnas – Outi Kröger – Timo Rauhanen (toim.). Helsinki 2009.
56. Local public sector in transition: A Nordic perspective. Antti Moisio (Ed.). Helsinki 2010.
57. Encouragement and discouragement. Essays on taxation and government expenditure. Tuomas Kosonen. Helsinki 2011.
58. Three takes on sustainability. Juha Honkatukia (Ed.). Helsinki 2011.
59. Talouden rakenteet 2011. Aki Kangasharju – Outi Kröger – Timo Rauhanen (toim.). Helsinki 2011.



VALTION TALOUDELLINEN TUTKIMUSKESKUS
STATENS EKONOMISKA FORSKNINGSCENTRAL
GOVERNMENT INSTITUTE FOR ECONOMIC RESEARCH

Valtion taloudellinen tutkimuskeskus
Government Institute for Economic Research
P.O.Box 1279
FI-00101 Helsinki
Finland

ISBN 978-952-274-011-3
ISSN 0788-4990



9 789522 740113