

Essays in Labor and Health Economics - on the Effects of Job Loss, Gender Norms and Physician Strikes

Von der Mercator School of Management, Fakultät für Betriebswirtschaftslehre, der
Universität Duisburg-Essen
zur Erlangung des akademischen Grades
eines Doktors der Wirtschaftswissenschaft (Dr. rer. oec.)
genehmigte Dissertation

von

Lea Nassal

aus

Limburg a.d. Lahn

Referentin: Prof. Dr. Marie Paul
Korreferent: Prof. Dr. Philip Jung
Tag der mündlichen Prüfung: 12.03.2024

Acknowledgment

I thank my first advisor, Marie Paul, for her guidance, and support. I am deeply thankful for her patience and dedication to helping me navigate through my PhD journey. I am equally indebted to my second supervisor, Philip Jung, for his valuable feedback.

I am very grateful to my co-authors on the second and third chapter of this dissertation, including Daniel Avdic, Seema Jayachandran, Martin Karlsson, Matthew Notowidigdo, Marie Paul, Heather Sarsons, Nina Schwarz and Elin Sundberg for the invaluable and enjoyable collaboration that significantly expanded the scope of this dissertation.

Many thanks for helpful comments and support to all members of the Research Training Group (RTG) Regional Disparities & Economic Policy, and my colleagues in Duisburg. Special thanks go to Karolin Süß, Nuan Stahl, Max Schäfer, Fabian Dehos, Philipp Markus, Malte Borghorst, and Lu Wei, for all their helpful comments and support throughout the years.

I thank the Research Data Center of the Institute for Employment Research (IAB) for generously providing the data and the German Research Foundation (DFG) via the RTG Regional Disparities & Economic Policy for financial support.

Finally, this journey would not have been possible without the encouragement and support of my family, to whom I'm very grateful.

Contents

Introduction	1
Bibliography	4
Main Chapters	7
1 Job Loss and Retirement	7
1.1 Introduction	8
1.2 Institutional Setting	12
1.3 Data	14
1.3.1 German Administrative Data	14
1.3.2 Outcome Variables	15
1.3.3 Plant Closures	16
1.3.4 Analysis Sample	17
1.4 Empirical Strategy	18
1.4.1 Matching Procedure	18
1.4.2 Descriptives	20
1.4.3 Empirical Specifications	24
1.5 The Effect of Job Loss on Retirement	26
1.5.1 Overall Effect	26
1.5.2 Heterogeneity Analysis	30
1.5.3 The Lifetime Costs of Job Loss	40
1.5.4 Robustness Analysis	43
1.5.5 Comparison with Effect of Late-Career Job Loss	46
1.6 Potential Channels	49
1.6.1 Pathways into Retirement	49
1.6.2 Alternative Channels	51
1.7 Conclusion	53
Bibliography	54

Appendix	59
1.A Institutional Setting	59
1.A.1 Pension Benefit Calculation	59
1.A.2 Pathways	60
1.A.3 Unemployment Insurance	61
1.B Data	61
1.B.1 Data Preparation	61
1.B.2 Approximation of Pension Benefits	63
1.C Appendix Figures and Tables	66
2 Moving to Opportunity - Together	77
2.1 Introduction	78
2.2 Data	82
2.2.1 German Data	82
2.2.2 Swedish Data	83
2.2.3 Moving Across Commuting Zones	84
2.2.4 Sample Selection, Variable Definition, and Descriptive Statistics	85
2.3 Empirical Strategy	88
2.4 Results	89
2.4.1 Descriptive Results	89
2.4.2 Main Results: Earnings Effects of Moving Across Com- muting Zones	92
2.4.3 Heterogeneity	96
2.4.4 Mass Layoff Results	98
2.5 Model-Based Estimation	99
2.5.1 Model	100
2.5.2 Heterogeneity in the Effects of Migration on Earnings by Female Share of Household Income	104
2.5.3 Model-Based Estimation	108
2.5.4 Additional Implications of $\beta < 1$: Gender Differences in Effects of Job Layoffs on Relocation and Gender Differ- ences in Child Penalties	112
2.6 Additional Evidence and Alternative Explanations	113
2.6.1 Gender Norms: East and West Germany	114

2.6.2	Alternative Explanations	116
2.7	Conclusion	118
	Bibliography	121
	Appendix	125
2.A	Proofs of Theoretical Results in Main Text	125
2.A.1	Additional Theoretical Results	131
2.A.2	Model Extensions	136
2.A.3	Model-Based Simulations	137
2.A.4	Extended Model of Child Penalty	137
2.B	Appendix Figures and Tables	139
2.B.1	Other Employment Measures	139
2.B.2	Heterogeneity	142
2.B.3	Predicted Income Methodology and Results	146
2.B.4	Descriptive Figures for Layoffs	149

3 Alive and Kicking? Short-Term Health Effects of a Physician Strike in Germany **153**

3.1	Introduction	154
3.2	Background	157
3.2.1	The German Hospital Sector	157
3.2.2	Marburger Bund	157
3.2.3	Physician Strike	159
3.3	Data	159
3.3.1	Strike Data	159
3.3.2	Hospital Level Data	160
3.3.3	Mortality Census	162
3.4	Econometric Strategy	163
3.4.1	Empirical Model	164
3.4.2	Patient Selection	166
3.4.3	Analysis of Spillover Effects	167
3.5	Results	167
3.5.1	Hospital Admissions and Patient Mortality	167
3.5.2	Care Rationing and Hospital Spillovers	171
3.5.3	Robustness Checks	174
3.6	Conclusion	178

Bibliography	180
Appendix	183
3.A Appendix Figures and Tables	183
3.B Variable Definitions	187

Concluding Remarks	189
---------------------------	------------

List of Figures

1.1	Displaced and Matched Control Workers before and after Displacement - Raw Means after Matching	23
1.2	Effect of Job Loss on Retirement - Overall Effect	29
1.3	Heterogeneous Effects by Age at Job Loss	31
1.4	Heterogeneous Effects by Pre-displacement Wage	33
1.5	Heterogeneous Effects by Economic Conditions at Job Loss	35
1.6	Heterogeneous Effects by Pre-displacement Education	37
1.7	Heterogeneous Effects by Pre-displacement Occupation	39
1.8	Effect of Job Loss on Retirement - Late Career Job Loss	48
1.9	Pathways into Retirement	50
1.10	Effect of Job Loss on Number of Sick Leave Days	52
1.C.1	Descriptives Data Exit (Raw Means) - Displaced and Matched Control Workers	70
1.C.2	Descriptives Data Exit (Raw Means) - Displaced and Matched Control Workers (by Age at Job Loss)	71
1.C.3	Displaced and Control Workers before and after Displacement - Raw Means before Matching	72
1.C.4	Pathways into Retirement (Raw Means) - Displaced and Matched Control Workers	73
1.C.5	Pathways into Retirement - By Age at Job Loss	74
1.C.6	Effect of Job Loss on Retirement - Long-run	76
2.1	Maps of Commuting Zones	90
2.2	Relationship between Moving and Labor Earnings and Employment	91
2.3	Impact of Move on Labor Earnings and Employment	94
2.4	Proportional Impact of Move on Wage Income	95
2.5	Impact of Move on Wage Income – By Age Groups	97

2.6	Impact of Move on Wage Income – By Gender-blind Predicted Female Share of HH Income	106
2.7	East vs. West German Origin	115
2.8	East vs. West German Origin, Men	115
2.9	Wage Income Results by Occupation - Below 50%, Germany	117
2.10	Wage Income Results by Occupation - Above 50%, Germany	117
2.11	Distance as outcome (in km) - Sweden	118
2.B.1	Relationship between Moving and Other Employment Measures	140
2.B.2	Event Study Results on Other Measures of Employment . .	141
2.B.3	Impacts of Move on Wage Income - By Timing of First Joint Child, Sweden	143
2.B.4	Impact of Move on Wage Income – By Gender-specific Predicted Female Share of HH Income	144
2.B.5	Impact of Move on Wage Income – By Predicted Female Share of HH Income, Median	145
2.B.6	Predicted Wage Income, Movers	147
2.B.7	Predicted Female Share of HH Income, Movers	148
2.B.8	Relationship between Layoffs and Labor Earnings and Employment	150
3.1	Location of Treatment and Control Hospitals	160
3.2	Event Study Results - Hospital Admissions	168
3.3	Event Study Results - Hospital Mortality	169
3.4	Event Study Results - District Mortality	171
3.A.1	Age Distribution by Emergency Status	184
3.A.2	Descriptives Mortality Rate	185
3.A.3	Descriptives Cases	185
3.A.4	Alternative Estimator - Mortality Rate	186
3.A.5	Alternative Estimator - Cases	186

List of Tables

1.1	Descriptive Statistics of Displaced and Non-Displaced Workers	21
1.2	Effect of Job Loss on Retirement - Overall Effect	28
1.3	Life-time Costs of Job Loss	41
1.4	Effect of Job Loss on Retirement - Robustness Checks	44
1.A.1	Requirements under Different Pathways	61
1.A.2	Statutory Retirement Ages under Different Pathways	62
1.A.3	Maximum UI Benefit Duration at Different Ages (in Months)	63
1.B.4	List of Occupations Excluded from the Analysis	64
1.C.5	Occupations of Displaced and Non-Displaced Workers	66
1.C.6	Lifetime Costs of Job Loss - Further Subgroups	67
1.C.7	Effect of Job Loss on Retirement Probability - Late Career Job Loss	69
1.C.8	Effect of Job Loss on Retirement - Women	69
2.1	Summary Statistics for Movers Sample	86
2.2	Summary Statistics for Job Layoffs Sample	87
2.3	Impact of Layoffs on Moving Probability	99
2.4	How Do the Effects of Moving by Gender Vary with the Predicted Female Share of Household Income?	107
2.5	Model Parameter Estimates	109
2.6	Assessing Model Fit	111
2.7	Model-Based Simulations	113
2.B.1	Stepwise Restrictions to Layoffs Sample, Sweden	151
3.1	Chronology of Events	158
3.2	Descriptive Statistics - Hospital Data	161
3.3	Descriptive Statistics - Mortality Data	163
3.4	Covariate Balancing Tests	165
3.5	Difference-in-differences Estimates - Hospital Admissions	169

3.6	Difference-in-differences Estimates - Hospital Mortality . . .	170
3.7	Difference-in-differences estimates: Patient composition . . .	172
3.8	Difference-in-differences Estimates - Heterogeneity by Emer- gency Status	173
3.9	Difference-in-differences Estimates - Spillover Analysis . . .	174
3.10	Difference-in-differences estimates: Placebo Regressions for 2004	175
3.11	Difference-in-differences Estimates - Alternative Control Group	176
3.12	Difference-in-differences Estimates - Alternative Treatment Indicator	177
3.13	Pre-trend Test	178
3.A.1	Physicians and Health Care Workers in Hospitals	183
3.A.2	Unweighted Mortality Regressions	183
3.A.3	Admissions by Cause	184

Introduction

Important events, such as relocation, job losses, or strikes, can have large effects on individuals or households. This dissertation seeks to analyze the effects of three distinct events: job loss, moves and strikes. The first chapter examines the long-term impact of job loss on retirement. The second chapter studies the impact of gender norms in couples' joint location decisions, and the last chapter the effects of a physician strike.

In chapter 1, I examine the long-term effects of job loss on retirement. Job loss is a large negative career shock that affects many workers – even years after the layoff (e.g. Jacobson et al., 1993; Couch and Placzek, 2010; Schmieder et al., 2023). Yet, we still know little about the lifetime costs of job loss and its interaction with workers' retirement decisions. In chapter 1, I provide first evidence on the long-term effects of job loss on retirement.

Exploiting plant closures, I compare the retirement behavior of displaced workers with similar workers who did not experience a layoff. I show that displaced workers adjust their retirement behavior by delaying retirement. Despite the adjustment, they still experience large losses in estimated pension benefits and lifetime income. Overall, displaced workers experience losses in estimated pension benefits of about 5% and in the present discounted value (PDV) of income of around 16%. Using non-displaced “twin” workers as counterfactuals, I show, that displaced workers would have experienced much larger losses if they did not delay retirement.

Chapter 1 improves our understanding of the long-term effects of job loss. As such it contributes to the large literature on the effects of job loss (e.g. Jacobson et al., 1993; Couch and Placzek, 2010; Schmieder et al., 2023; Davis and von Wachter, 2011). While some papers have examined the effects of late-career job loss (Chan and Huff Stevens, 2001; Chan and Stevens, 2004; Merkurieva, 2019), this paper provides new evidence on the long-term effects of job loss on retirement, focusing on job loss among young and middle-aged workers.

In chapter 2 (co-authored with Seema Jayachandran, Matthew Notowidigdo, Marie Paul, Heather Sarsons and Elin Sundberg), we investigate the effects of relocation on men's and women's earnings, and test how much of the gender

earnings gap from relocation is due to differences in earnings potential versus gender norms. Early models of the household predict that couples will make location decisions to maximize joint income (Mincer, 1978; Frank, 1978). Joint location decisions may therefore result in a gender earnings gap if men have higher earnings or earnings potential than women. However, a gender norm that prioritizes men's career advancement may also explain the gender earnings gap.

Using administrative data from Germany and Sweden, we first show that moves disproportionately benefit men. Men's earnings increase sharply over the first five years following the move, whereas women experience only a modest earnings increase. These results may be due to men's higher earnings or gender norms. To distinguish between these different potential explanations, we consider a model of household decision-making in which households potentially place more weight on income earned by the man relative to the woman. In a standard unitary model of joint income maximization, moves should not systematically benefit men in couples where the woman and the man have identical pre-move earnings and earnings potential. Further, the earnings gender gap after relocation is smaller the higher the woman's share of household income.

Our results show that the post-move earnings gap is indeed smaller among couples in which the woman has a higher predicted share of household income. But even among couples where women have higher potential earnings, we find that men benefit more from the move than women. Using our model, we then test and reject the unitary model in both countries, with larger deviations from the unitary household benchmark in Germany than in Sweden. Households in both countries place less weight on income earned by the woman compared to the man, particularly in Germany.

This chapter contributes to the large literature studying sources of gender gaps in labor market outcomes (e.g. Bertrand et al., 2015; Kleven et al., 2019; Cortes et al., 2021; Goldin and Katz, 2016). We add to this literature by studying an aspect, which has not received much attention: women may take less advantage of career enhancing long-distance moves or may even experience earnings losses as a "tied mover".

In chapter 3 (co-authored with Daniel Avdic, Martin Karlsson and Nina Schwarz), we study the effects of a nationwide physician strike in Germany.

Physician strikes are often seen as controversial, as disruptions in access to healthcare providers may have large negative impacts on the quality of healthcare. A priori, the extent of the adverse effects is not clear. On the one hand, strikes may affect the quality of healthcare due to reduced geographical access to providers, higher workload among non-striking staff and crowding effects (Avdic, 2016; Piérard, 2014; Lin, 2014; Aiken et al., 2002). On the other hand, the adverse effects of strikes can be mitigated by rationing care and triaging patients based on urgency. We provide evidence on the extent of these adverse effects by analyzing a physician strike, which took place in 2006.

Drawing on rich administrative data on all hospital admissions and deaths from 2000-2008, we compare outcomes of striking and non-striking hospitals in a difference-in-differences model. Our results show that the strike led to a large reduction in hospital admissions in striking hospitals, while in-hospital mortality and emergency admissions increased. However, once we adjust for changes in the patient composition during the strike, the mortality effect is reduced by half. Further, we find some spillover effects to nearby hospitals, whereas there is no post-strike catch-up effect on admissions, and no effect on mortality outside of hospitals. Overall, our results suggest that healthier patients avoided care or were triaged, resulting to a higher share of patients with higher underlying mortality risk in striking hospitals.

This chapter adds to the literature on strikes in health care facilities. While some papers have studied the effects of healthcare workers (Gruber and Kleiner, 2012; Kronborg et al., 2016; Friedman and Keats, 2019; Hirani et al., 2019), less is known about strikes of physicians. To our knowledge, only two papers have investigated the impact of physician strikes (Stoye and Warner, 2023; Costa, 2019). In comparison to these studies, we examine the impact of a longer strike that affected entire hospitals, providing a very powerful test of the system's resilience in the event of a major unexpected challenge.

In sum, each chapter of this dissertation leverages large administrative datasets to provide new perspectives on important topics, relevant for both policymakers and the scholarly community.

Bibliography

- Aiken, L. H., Clarke, S. P., Sloane, D. M., Sochalski, J., and Silber, J. H. (2002). Hospital nurse staffing and patient mortality, nurse burnout, and job dissatisfaction. *Jama*, 288(16):1987–1993.
- Avdic, D. (2016). Improving efficiency or impairing access? health care consolidation and quality of care: Evidence from emergency hospital closures in sweden. *Journal of health economics*, 48:44–60.
- Bertrand, M., Kamenica, E., and Pan, J. (2015). Gender Identity and Relative Income within Households *. *The Quarterly Journal of Economics*, 130(2):571–614.
- Chan, S. and Huff Stevens, A. (2001). Job loss and employment patterns of older workers. *Journal of Labor Economics*, 19(2):484–521.
- Chan, S. and Stevens, A. H. (2004). How does job loss affect the timing of retirement? *Contributions in Economic Analysis Policy*, 3(1):1–24.
- Cortes, P., Pan, J., Pilosoph, L., and Zafar, B. (2021). Gender differences in job search and the earnings gap: Evidence from business majors. *National Bureau of Economic Research Working Paper*, (28820).
- Costa, E. (2019). License to kill? the impact of hospital strikes. *The Impact of Hospital Strikes (May 20, 2019)*.
- Couch, K. A. and Placzek, D. W. (2010). Earnings losses of displaced workers revisited. *American Economic Review*, 100(1):572–89.
- Davis, S. J. and von Wachter, T. (2011). Recessions and the costs of job loss. *Brookings Papers on Economic Activity*, pages 1–73.
- Frank, R. H. (1978). Why women earn less: The theory and estimation of differential overqualification. *The American Economic Review*, 68(3):360–373.
- Friedman, W. and Keats, A. (2019). Disruptions to health care quality and early child health outcomes: Evidence from health worker strikes in kenya. mimeo.

- Goldin, C. and Katz, L. F. (2016). A most egalitarian profession: Pharmacy and the evolution of a family-friendly occupation. *Journal of Labor Economics*, 34(3):705–746.
- Gruber, J. and Kleiner, S. A. (2012). Do strikes kill? evidence from new york state. *American Economic Journal: Economic Policy*, 4(1):127–57.
- Hirani, J., Sievertsen, H. H., and Wüst, M. (2019). Beyond treatment exposure: The timing of early interventions and children’s health. mimeo.
- Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993). Earnings losses of displaced workers. *The American economic review*, pages 685–709.
- Kleven, H., Landais, C., and Søgaaard, J. E. (2019). Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics*, 11(4):181–209.
- Kronborg, H., Sievertsen, H. H., and Wüst, M. (2016). Care around birth, infant and mother health and maternal health investments—evidence from a nurse strike. *Social Science & Medicine*, 150:201–211.
- Lin, H. (2014). Revisiting the relationship between nurse staffing and quality of care in nursing homes: An instrumental variables approach. *Journal of health economics*, 37:13–24.
- Merkurieva, I. (2019). Late career job loss and the decision to retire. *International Economic Review*, 60(1):259–282.
- Mincer, J. (1978). Family migration decisions. *Journal of Political Economy*, 86(5):749–773.
- Piérard, E. (2014). The effect of physician supply on health status: Canadian evidence. *Health Policy*, 118(1):56–65.
- Schmieder, J. F., Von Wachter, T., and Heining, J. (2023). The costs of job displacement over the business cycle and its sources: evidence from germany. *American Economic Review*, 113(5):1208–1254.
- Stoye, G. and Warner, M. (2023). The effects of doctor strikes on patient outcomes: Evidence from the english nhs. *Journal of Economic Behavior & Organization*, 212:689–707.

Job Loss and Retirement

Chapter Abstract

This paper provides new evidence on the long-term effects of job loss on retirement, based on German administrative data from 1975 to 2021. To identify the effect of job loss, I exploit plant closures to compare the retirement behavior of displaced workers with similar workers who did not experience job loss. I show that displaced workers delay their retirement in response to the shock. However, even if displaced workers adjust their retirement behavior, they still experience losses in estimated monthly pension benefits. Overall, the lifetime costs of job loss are large: displaced workers experience losses in the present discounted value (PDV) of income of around 16%.

I thank the German Institute for Employment Research (RDC-IAB) for generously providing the data and support with running programs remotely.

1.1 Introduction

Job loss is a persistent negative career shock that affects many workers. In the US, for example, approximately 8.6 million workers involuntarily lost their jobs from 2019 to 2021 (US Bureau of Labor Statistics, 2022)¹. Given this relevance, the high costs of job loss are well-documented in the literature. Displaced workers experience significant earnings losses, even years after their layoff (e.g. Jacobson et al., 1993; Couch and Placzek, 2010; Schmieder et al., 2023). Yet, we still know little about the lifetime costs of job loss and whether it affects workers' retirement decisions². This is striking, as job loss may not only lead to large earnings losses, but also to reduced future pension benefits. For some workers, this may increase the risk of old-age poverty, which is a major problem given increasing life expectancy and reduced pension replacement rates. Key questions are therefore whether displaced workers adjust their retirement behavior in response to the shock and how large the lifetime costs of job loss are.

Despite their importance, these questions remain understudied. This is largely due to little comprehensive data that allow to track workers from displacement until retirement. Some work examines how late-career job losses influence older workers' retirement decisions (Chan and Huff Stevens, 2001; Chan and Stevens, 2004; Merkurieva, 2019), but workers' response to the job loss likely varies depending on when the shock occurs. It is still unresolved how job loss among young and middle-aged workers affects their future retirement decisions and how large the lifetime costs of job loss are.

In this paper, I provide first evidence on the *long-term* effects of job loss on retirement. Drawing on more than 40 years of German administrative data, I show that displaced workers adjust their retirement behavior by delaying retirement. Despite the adjustment, they still experience large losses in estimated pension benefits and lifetime income. Overall, displaced workers experience losses in the PDV of income of around 16%. These results demonstrate the importance of connecting the job loss literature to what

¹3.6 million of the 8.6 million were workers with more than three years of tenure (US Bureau of Labor Statistics, 2022).

²Throughout the paper, I use the term retirement for simplicity. I define the year of retirement entry as the year in which a worker permanently exits the labor market (more information in section 1.3.2).

happens to workers' retirement decisions and improve our understanding of the lifetime costs of job loss.

For the empirical analysis, I use German administrative data, with several advantages. The data includes detailed information on workers' labor market biographies spanning more than 40 years, which I use to identify plant closures and track workers from job loss until retirement. In addition, the data includes a rich set of personal and job-related characteristics. This enables me to control for characteristics that could influence workers' retirement behavior, such as occupation, education, or birth cohort.

To establish a causal link between job loss and retirement, I exploit a quasi-exogenous setting in which high-tenured male workers lose their jobs due to plant closures. Displacement is likely to be unexpected and costly for these workers, as they would most likely not have changed their long-term jobs if there was no closure. However, even in this setting, displaced workers may retire differently than non-displaced workers, simply because of different worker characteristics. To compare displaced workers with non-displaced workers, I match each displaced worker with a non-displaced "twin", who has similar observable characteristics before the job loss. This allows me to flexibly compare the retirement behavior of displaced workers with otherwise similar workers who did not experience job loss.

I show that displaced workers react to their job loss by delaying retirement. Displaced workers are about 10% less likely to be retired between the ages of 60 and 64 compared to similar non-displaced workers. Even if displaced workers adjust their retirement behavior, they still experience losses in estimated pension benefits of about 5%. Losses in estimated pension benefits would however be much larger had displaced workers not adjusted their retirement behavior. Using non-displaced "twin" workers as counterfactuals, I show that displaced workers would have experienced losses in estimated pension benefits of about 11% if they did not delay retirement.

Having identified the overall effect and shown the existence of important responses in workers' retirement behavior, I document interesting effect heterogeneity across different sub-groups. First, I show that the delay in retirement entry differs by workers' age at job loss, adding to the evidence on life-cycle differences in the impact of job loss (Salvanes et al., 2022; Rinz, 2022). I show that the retirement probability from age 60-64 is about twice

as large for worker in the 35-40 and 41-45 age groups than in the 46-50 one. Older workers (46-50) are unemployed for a long time after losing their job, suggesting that they face difficulties in finding a new job (Farber et al., 2019; Carlsson and Eriksson, 2019) and leads some of them to retire early. Losses in estimated pension benefits are smaller for younger workers, whose adjustment in the retirement timing is the largest.

Second, some studies have shown that the costs of job loss vary by the prevailing economic conditions (Schmieder et al., 2023; Davis and von Wachter, 2011; Farber, 2017). For example, Schmieder et al. (2023) estimate that the costs of job loss nearly double in size for workers who are displaced during recessions. I show the delay in workers' retirement is much larger for workers displaced during a recession. Workers who were displaced during a recession are more than half as likely to be retired between the ages of 60 and 64 as workers who are not, showing that the response is larger for those who face a larger negative income shock. Finally, my results show that especially low-skilled workers delay retirement, whereas losses in estimated pension benefits are larger for high-skilled workers and workers in high-paying occupations such as mechanics and engineering.

The losses in estimated pension benefits increase the lifetime cost of job loss. Overall, displaced workers experience losses in the PDV of income (including estimated pension benefits) of around 16%. To put this magnitude into context, Davis and von Wachter (2011) estimate the loss in the PDV of earnings for the US to be around 12%, considering only a 20-year period after displacement. This shows the importance of going beyond the time frame previously studied to fully capture the lifetime costs of job loss. Losses differ by subgroups and are particularly large for older and high-skilled workers, as well as for workers with high pre-displacement wages – consistent with the previous findings on estimated pension benefits. Workers with high pre-displacement wages, for example, experience losses in the PDV of income which are about twice as large as for workers with low pre-displacement wages.

Taken together, my results show that job loss has long-lasting negative effects, even beyond working life. In response to the shock, displaced workers delay retirement but still experience losses in estimated pension benefits. For fairness considerations, this highlights the importance of policies that help

to quickly reintegrate displaced workers into the labor market to buffer the long-term cost of job loss.

The results of this paper make several contributions to the literature. First, the findings improve our understanding of the long-term effects of job loss. As such, this paper relates to a large body of work, going back to at least Jacobson et al. (1993), that examines the impacts of job loss. Prior research has shown that job loss leads to large and long-lasting earnings losses (e.g. Couch and Placzek, 2010; Davis and von Wachter, 2011; Lachowska et al., 2020; Jacobson et al., 1993; Schmieder et al., 2023), increases job instability (Jarosch, 2023) and raises the incidence of future job losses (Stevens, 1997, 2001). For example, Schmieder et al. (2023) use German data and document that even 10 years after displacement affected workers experience earnings losses about 15%. By using administrative data that covers more than four decades, I contribute to this literature by providing evidence on the lifetime costs of job loss, including losses due to lower estimated pension benefits.

More closely related, some studies have examined how late-career job loss affects workers' retirement decisions. These studies document that late-career job loss leads to substantially lower re-employment probabilities (Chan and Huff Stevens, 2001) and earlier retirement (Chan and Stevens, 2004; Merkurieva, 2019). Coile and Levine (2007) and Card et al. (2014) find that workers claim pension benefits at the earliest possible age when faced with a negative labor market shock toward the end of their careers. Chan and Stevens (2004) show that this behavior cannot be explained by financial considerations, but rather barriers to re-employment (Farber et al., 2019; Carlsson and Eriksson, 2019) may play a more important role. Relative to these studies, I look into the long-term impact of job loss and document that young and middle-aged workers respond to the job loss by delaying retirement. This shows that workers' response varies depending on the age at which they experience the shock, complementing evidence from Salvanes et al. (2019) and Rinz (2022), who document different impacts of job loss across the life-cycle.

Finally, this paper also relates to the literature about workers' retirement behavior. The retirement decision is one of the most well-studied decisions in economic research and a number of theoretical models of retirement decisions focus on different approaches to specify the financial incentive to

retire (Stock and Wise, 1990; Gustman et al., 1986; Rust and Phelan, 1997). More recently, a large number of studies took advantage of pension reforms to investigate how changes in statutory retirement ages affect workers' timing of retirement (Brown, 2013; Manoli and Weber, 2016; Seibold, 2021), as well as how such amendments lead to program substitution effects (Inderbitzin et al., 2016; Geyer and Welteke, 2021). Whereas those studies exploit quasi-experimental variation in statutory retirement ages to show the short-term effects on workers' retirement behavior, I examine how a negative labor market shock influences workers' retirement behavior in the long term.

1.2 Institutional Setting

The following section describes the main features of the German public pension system and addresses the most relevant features of the unemployment insurance (UI) system and its interaction with the pension system.

Key Features of the Public Pension System. The German public pension system has a pay-as-you-go scheme and covers most private sector employees³. For most retirees, public pension is the most important income source, with income of occupational pensions or individual retirement accounts do not play a major role (Börsch-Supan and Wilke, 2004). Pension benefits are calculated according to a pension formula based on workers lifetime contribution history⁴. In Germany, all pension contributions collected throughout life are taken into account to calculate the pension benefits. Pension contributions cannot only be acquired during periods of employment, but also during other insurance periods such as unemployment, sickness, military service, or child raising. The number of pension contribution points is roughly proportional to workers' earnings or replacement payments, but there is a maximum number of points a worker can contribute per year. In general, workers with only a few contribution points will receive a low pension.

There are two types of statutory retirement ages in the pension system: the *normal retirement age (NRA)* and the *early retirement age (ERA)*, which

³Most self-employed people are exempted from participating in the public pension system and civil servants have a separate pension system. In the analysis, civil servants and self-employed people are not considered.

⁴There is more information on the calculation of pension benefits in appendix 1.A.

both depend on workers' birth cohort, gender, and contribution history. The NRA is the age at which a worker can claim full pension. For workers in the analysis sample whose birth cohort is between 1943 and 1952, the NRA varies between 60 and 65. In contrast, the ERA is the age at which a worker can claim pension at the earliest point, but only with deductions. Specifically, a deduction of 0.3% is imposed for each month a worker retires before reaching the NRA. If a worker decides to go into early retirement two years before reaching the normal retirement age, the pension benefits would be 7.2% less. For workers in the analysis sample, the ERA varies between 60 and 63. In addition to retiring earlier, it is also possible for workers to work beyond the NRA, if an employer agrees to extend their contract. For each month a worker retires after the NRA, an additional reward of 0.5% is paid. Displaced workers can therefore react to job loss by choosing to retire earlier or later. Depending on workers' pathway into retirement, they have to fulfill different requirements to claim a pension⁵. For claiming a regular pension at age 65, the only requirement is to have had at least five contribution years (Rentenversicherung, 2020).

Marginal Employment while Receiving Pension Benefits. Pensioners can also decide to work while receiving pension benefits. They do however face a strict earnings test between the ERA and NRA, where earnings above 450 € per month lead to reduced pension benefits (Ye, 2022). Early retirement while working in a regular job is therefore not attractive. For some workers it may however be attractive to work in marginal employment while receiving pension benefits, to supplement a low pension. The most popular type of marginal employment in Germany is the so-called "*mini-job*", which is exempted from social security contributions, income taxation, and participation in the pension system⁶. In those jobs, workers can earn a maximum of 450 € per month, which does not lead to any reduction in pension benefits⁷.

⁵There is more information on the different pathways in appendix 1.A.

⁶Until 2012, workers in mini-jobs were exempted from participating in the public pension system. They did not have to make any pension contributions. They could however make voluntary contributions so that they could get full pension insurance coverage. Since 2013, they have had to make pension contributions and get full pension insurance coverage. Upon request, they can still get exempted from paying contributions, but then they no longer have pension insurance coverage (Deutsche Rentenversicherung, 2014).

⁷The mini-job earnings threshold was 325 € from 1999 to 2003, 400 € from 2004 to 2013, 450 € from 2013 to 2022 and is 520 € from 2023 onward (Gudgeon and Trenkle, ming;

UI Benefits as Bridge to Retirement. Many workers do not directly transition from employment to retirement but use UI as a stepping stone into retirement (Hairault et al., 2010; Gudgeon et al., 2019; Giesecke and Kind, 2013; Inderbitzin et al., 2016). In Germany, unemployed workers receive about 60% of their last net income as replacement payments and job search requirements are very low for older workers. Depending on workers' birth cohort and age, the maximum UI benefit duration ranged from 18 to 32 months for workers aged 55 or above⁸. Workers acquire pension contributions during periods of receiving UI benefits, which increase future pension benefits and make it attractive to use unemployment as a stepping stone into retirement.

1.3 Data

1.3.1 German Administrative Data

For the empirical analysis, I use a sample of employment biographies from 1975 to 2021. In particular, I use a 2% random sample of all workers subject to social security contributions provided by the *Institute for Employment Research (IAB)*⁹.

Three features of the data make it suitable for studying the long-term effects of job loss on retirement. First, the data allows me to observe a large number of workers and their labor market biographies over a long period. The data covers more than four decades and I can follow each worker from the time of displacement until retirement. Second, for each employee, the data contains information on the employer that enables me to link plant closures to workers. Finally, the data includes a rich set of personal and job-related characteristics, enabling me to compare displaced workers with similar non-displaced workers. A caveat is that the data does not include information on pension claims and benefits. I therefore use the year of labor market exit as a proxy for the year of retirement entry (more detail in section

Deutsche Rentenversicherung, 2014).

⁸There is more detail in appendix A.3.

⁹This study uses the Weakly anonymous version of the Sample of Integrated Labour Market Biographies (SIAB) – Version 7521 v1 (DOI: 10.5164/IAB.SIAB7521.de.en.v1). This data does not include civil servants and self-employed people. Civil servants benefit from extensive employment protection and are unlikely to matter for studying the effects of job loss. Civil servants and most self-employed people also do not participate in the public pension system.

1.3.2).

The data consists of day-to-day information on all periods in employment covered by social security, all periods of receiving UI benefits, and all periods registered as searching for a job. Each period contains information on the corresponding wages and benefit levels¹⁰. The wage information is accurate, as the employer has to report wages for social security purposes. Like most social security data, wages are however right-censored at the social security contribution ceiling. I impute right-censored wages using a two-step imputation procedure following Dustmann et al. (2009) and Card et al. (2013). The data also includes personal and job-related characteristics such as gender, education, occupation, nationality, and year of birth. For each employee, the data contains information on the employer such as number of employees, location and industry.

1.3.2 Outcome Variables

Retirement. Throughout the analysis, I consider the *probability to be retired* as a main outcome variable. Retirement is defined as an absorbing state and the year of retirement is defined as the year in which a worker is last observed in the data as either employed or unemployed. I allow workers to work in small, irregular employment relationships while in retirement. More specifically, I consider workers as retired if they only work in marginal employment in a given year and have already reached the earliest possible retirement age. Note that some workers may use unemployment as a stepping stone into retirement and therefore the year of retirement and the year of permanent employment exit may differ.

Estimated Pension Benefits. In addition to retirement, I consider *estimated pension benefits* as a second outcome variable. Unfortunately, the data does not directly include information on pension benefits, therefore I approximate pension benefits using information on workers' labor market biographies. In the data, I have information on all periods of employment and all periods receiving UI benefits, which are the main determinants of pension benefits. I calculate each worker's estimated pension benefits using the pension for-

¹⁰The data includes information on the level of UI I benefits, but not on level UI II benefits (unemployment assistance).

mula¹¹. This formula takes into account workers' full contribution history, meaning that pension benefits depend on workers' lifetime contribution history and not only on the last or best years, as in some other countries. In addition, it accounts for workers' retirement age by penalizing early retirement and rewarding retirement after normal retirement age.

Additional Outcomes. To investigate workers' pathways into retirement, I also consider the following outcome variables: *employment*, *unemployment* and *inactivity*. All these outcomes are defined as non-absorbing states. A worker is considered as employed if he is mainly working in a given year and unemployed if he is mainly unemployed (receiving UI benefits). A worker is defined as inactive if he is temporarily not in the labor market in a given year. Note that, by construction, workers cannot enter retirement after being inactive – they can switch between employment, unemployment and inactivity, as long as they are not yet retired.

1.3.3 Plant Closures

To identify plant closures, I use an extension file that contains information on the type of plant closures. This helps to distinguish “true” closures from those that are merely spin-offs of existing plants, takeovers, or ID changes (Hethey-Maier and Schmieder, 2013). A plant is considered as having a *plant closure* if its operating ID disappears between June 30 in two consecutive years. To avoid identifying restructuring of plants instead of “true” closures, it is required that no more than 40% of the displaced workers are reemployed at the same workplace in the year after the displacement. They are then either unemployed or reemployed at different workplaces.

Since closures are identified by the percentage of workers who leave, only plants with at least 20 employees are considered. Throughout the analysis, I consider West German plants, as information on East Germany is only available from 1992 onwards.

¹¹There is more information on the calculation of estimated pension benefits in appendix 1.B.

1.3.4 Analysis Sample

In line with the literature, displaced workers are considered as workers who involuntarily separate from their long-term jobs due to an exogenous shock (Jacobson et al., 1993). For these workers, displacement is likely to be unexpected and costly, as they would most likely not have changed their jobs otherwise. To focus on these workers, I apply the following *sample restrictions*: in the year before the displacement ($t - 1$), the worker is male, employed full-time, 35 to 50 years old and has at least three years of tenure with the same main employer (the employer from whom the worker receives the highest wage in a given year). I focus on men in the main analysis, as they have a higher attachment to the labor market, which allows a more accurate identification of the retirement entry and approximation of pension benefits¹². Since the goal is to estimate the effect of job loss on retirement, I only consider workers whose birth cohort is between 1943 and 1954. This allows me to observe each worker from age 35 to 66.

I then define a worker as displaced between year $t - 1$ and t if he leaves the plant between year $t - 1$ and t and the plant experiences a closure either (i) between year $t - 1$ and t or (ii) between year t and $t + 1$. The latter case refers to workers leaving the plant within a year before the final shutdown. “Early leavers” are included in the sample, because workers who are employed at the time of the final shutdown are most likely a non-random sample of displaced workers (Schwerdt, 2011; Pfann and Hamermesh, 2001). For each worker, I only consider the first displacement, as future outcomes may be influenced by the first displacement. To construct a control group of non-displaced workers, each non-displaced worker is randomly assigned a “*placebo job loss*”.

A caveat of the data is that it only covers individuals working in jobs subject to social security or receiving UI benefits. In addition to retirement there are other reasons why people may leave the data, for example becoming self-employed, starting a civil service job, emigrating or dying. For those

¹²Note that periods of childcare are taken into account when calculating pension benefits. Unfortunately, the data only allows approximating the birth of the first child for women, as the approximation relies on mothers being observed in the data before they give birth. Therefore, I cannot approximate pension benefits for women. Appendix table 1.C.8 presents the results of the effect of job loss on retirement timing for women. The results are comparable to those of the main sample, but the coefficients are quite imprecise estimated, as only 329 matched worker pairs could be identified.

workers, the year of retirement entry would not be correctly identified. I therefore apply additional restrictions to avoid misclassifications. First, I exclude workers without German citizenship in order to avoid misclassifying them as retired when in fact they emigrated. Second, I exclude workers in occupations with a high share of self-employment to prevent workers being misclassified as retired when in fact they became self-employed¹³¹⁴. Finally, I exclude deceased workers. All restrictions are applied to displaced as well as non-displaced workers. Before matching, the final sample consists of 1,467 displaced and 11,006 non-displaced workers.

1.4 Empirical Strategy

My analysis focuses on estimating the causal effect of job loss on retirement. To estimate this effect, I would ideally randomly assign job losses to workers. As this is not feasible, I exploit plant closures as a *quasi-experimental setup* and use matching to find appropriate counterfactuals for the displaced workers. In this section, I first describe the matching procedure and provide descriptive evidence for the labor market outcomes before and after displacement. I then describe the empirical specifications and identifying assumption.

1.4.1 Matching Procedure

Even if plant closures result in plausibly exogenous events for displaced workers, there are still observable differences between displaced and non-displaced workers that may influence workers' retirement decision and make comparison difficult. The sample of displaced workers is a selected sample, as plant closures are concentrated in some occupations and some worker groups are more likely to be affected. Table 1.1 shows that displaced and non-displaced workers differ before matching. Displaced workers are on average

¹³I also exclude workers in the public sector, as only 0.10% of displaced workers work in the public sector while this applies to 7.11% of non-displaced workers. Note that I only exclude individuals who work in the public sector and in occupations with a high share of self-employment in year $t - 1$. Appendix 1.B contains a list of excluded occupations.

¹⁴Appendix figures 1.C.1 and 1.C.2 show that there is no difference in data exits directly after the job loss between displaced and matched control workers, with the exception of workers aged 47 to 50 at the time of displacement. Those workers are more likely to exit directly after the displacement.

about one year older than non-displaced workers. They have a slightly higher unemployment experience and lower earnings. Overall, displaced workers are somewhat negatively selected compared to their non-displaced counterparts.

To identify appropriate counterfactuals for the displaced workers, I apply *coarsened exact matching* (Iacus et al., 2011, 2012) and match each displaced worker with a non-displaced control worker (without replacement). The pool of potential control workers comprises all non-displaced workers who never experienced a plant closure and fulfill the same restrictions as displaced workers. Note that there is no restriction that non-displaced workers cannot change employers or become unemployed between year $t - 1$ and t . Applying coarsened exact matching, each displaced is matched to a non-displaced control worker. In particular, I match exactly on the worker's age (35-50), birth cohort (1943-1954) and the displacement year (1978-2004), and coarsely on the worker's occupation (8 groups)¹⁵, education (3 groups)¹⁶ (all measured in $t - 1$), cumulative pension contribution points (2 groups) (measured in $t - 2$) and log wages (2 groups) (measured in $t - 3$ and $t - 4$)¹⁷. Log wages are measured in $t - 3$ and $t - 4$ to avoid picking up any pre-displacement wage losses.

Matching exactly on birth cohort is important for the analysis, as the statutory retirement age varies from one cohort to another. Matching on occupation and education allows me to compare workers with similar jobs, while matching on cumulative contribution points ensures that each displaced worker is matched to a non-displaced “twin” with a similar contribution history before the job loss. Finally, matching exactly on age and displacement year secures that comparison between displaced and matched control workers is done at the same age and calendar year.

¹⁵I distinguish between the following broad occupations: (i) resource extraction and production, (ii) construction, (iii) mechanics, engineering & technicians, (iv) transportation, sales & service, (v) accounting, management & law, (vi) arts, (vii) health, education & social affairs, (viii) others.

¹⁶Education groups are defined as follows: (i) no (recognized) completed education, (ii) vocational training/high school diploma and (iii) University degree.

¹⁷If multiple potential control workers can be identified, for each displaced worker the non-displaced worker with the closest pre-displacement log wage is selected. This matching procedure allows 84.6% of the displaced workers to be successfully matched to a non-displaced “twin”.

1.4.2 Descriptives

Table 1.1 shows *descriptive statistics* for displaced and non-displaced workers before and after matching. The first two columns report descriptive statistics for displaced and non-displaced workers before matching and columns 3 and 4 report descriptive statistics after matching.

Table 1.1: Descriptive Statistics of Displaced and Non-Displaced Workers

	Before Matching		After Matching	
	Displaced (1)	Controls (2)	Displaced (3)	Controls (4)
Age (years)*	43.37 (4.65)	42.39 (4.54)	43.16 (4.66)	43.16 (4.66)
Birth year*	1949.02 (3.45)	1948.33 (3.51)	1948.89 (3.45)	1948.89 (3.45)
Education				
No/unrecognised education*	0.10 (0.30)	0.10 (0.29)	0.05 (0.23)	0.05 (0.23)
Vocational training/high-school diploma*	0.84 (0.36)	0.80 (0.40)	0.90 (0.31)	0.90 (0.31)
University degree*	0.06 (0.24)	0.11 (0.31)	0.05 (0.22)	0.05 (0.22)
Job tenure at current plant (years)	11.06 (6.23)	11.18 (5.62)	10.93 (6.13)	12.04 (6.18)
Experience in employment (years) (censored in 1975)	15.79 (6.13)	14.39 (5.67)	15.75 (6.11)	15.86 (6.08)
Experience in unemployment (years) (censored in 1975)	0.27 (0.74)	0.18 (0.58)	0.25 (0.73)	0.22 (0.61)
Total yearly labor earnings	44962 (26860)	53716 (30952)	45360 (26531)	49559 (29683)
Days employed per year	342.10 (54.73)	359.35 (31.85)	342.58 (54.09)	360.48 (28.54)
Daily log wage	4.78 (0.39)	4.90 (0.41)	4.79 (0.39)	4.83 (0.37)
Daily log wage (t-3)*	4.78 (0.35)	4.89 (0.40)	4.79 (0.34)	4.81 (0.36)
Daily log wage (t-4)*	4.78 (0.35)	4.89 (0.40)	4.79 (0.34)	4.81 (0.36)
Cumulative contribution points*	19.02 (8.63)	18.09 (8.29)	19.20 (8.79)	19.19 (8.64)
Age retirement entry (censored after 66)	60.52 (5.50)	59.98 (5.67)	60.53 (5.55)	60.31 (5.28)
Observations	1467	11006	1220	1220

Notes: Statistics shown are means with standard deviations in parentheses. Characteristics are measured in pre-displacement year $t - 1$, if not stated differently. All workers fulfill the same baseline restrictions (see section 1.3.4). Characteristics with * are used in matching procedure. Right-censored observations contribute one observation at age 66 to the sample.

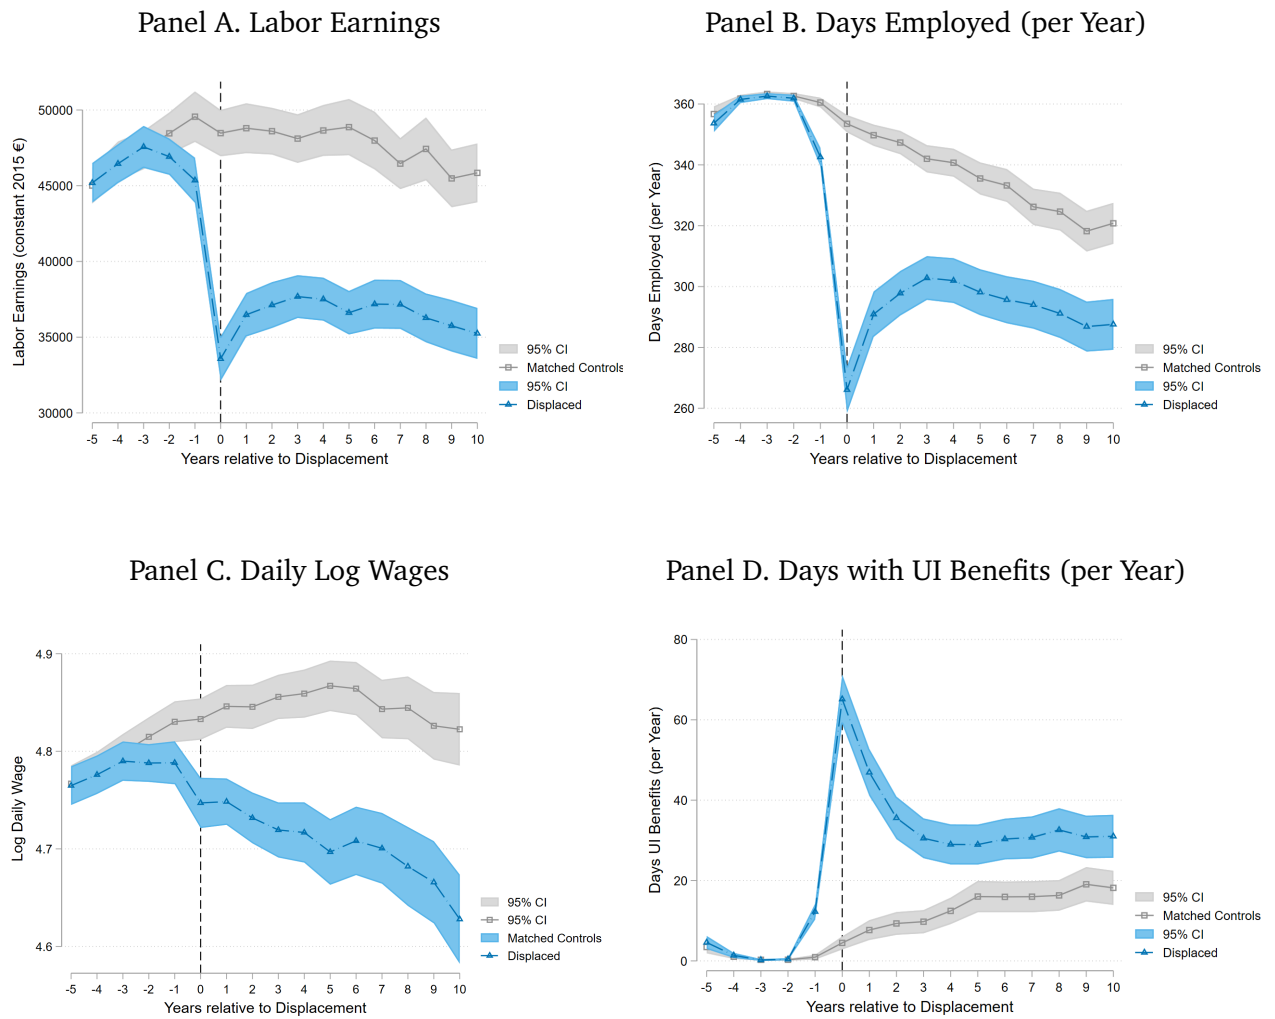
Before matching, displaced workers are somewhat negatively selected compared to non-displaced workers. Displaced workers have lower earnings and higher unemployment experience¹⁸. Compared to non-displaced workers, they enter retirement on average about six months later. After matching (columns 3 and 4), there are only small differences between displaced and non-displaced workers, indicating that the matching procedure works well. By construction, displaced and control workers are identical in terms of all coarsened covariates used in the matching. They are also more similar in terms of covariates not explicitly matched on. Displaced workers enter retirement about three months later than their matched control workers, as compared to about six months prior to matching.

Figure 1.1 shows the effect of job loss from five years before up to ten years after the job loss. In particular, it displays the means of labor market outcomes for displaced workers and their matched control workers. Figure 1.1 indicates that the pre-displacement trends of displaced and matched control workers are very similar, showing that the matched control group serves as an appropriate counterfactual¹⁹. Overall, the figure shows that job loss leads to long-term losses in earnings, employment and wages and significantly increases the receipt of UI benefits.

¹⁸Appendix table 1.C.5 shows workers' occupations before and after matching.

¹⁹Appendix figure 1.C.3 shows that even before matching, pre-displacement trends are very similar for displaced and non-displaced workers, but there are differences in levels.

Figure 1.1: Displaced and Matched Control Workers before and after Displacement - Raw Means after Matching



Notes: This figure shows raw means for displaced workers and their matched control workers. Plant closures take place between June 30th in $t - 1$ and June 30th in t . The outcome data are measured in calendar years and observed from 1975-2021. Workers are displaced in any year from 1978-2004.

Panel A shows that yearly earnings of displaced workers drop sharply in the year of displacement. Earnings are about 13,000 e (28%) lower in year 0 compared to the average before displacement. The earnings of displaced workers recover only partially in subsequent years and a significant earnings gap remains even ten years after the job loss.

Panel B and C show that earnings losses are explained by losses in employment and wages. Panel B indicates that displaced workers have about 100 fewer employed days (26%) in year 0 relative to the years before displacement. In the long term, employment recovers only partially and displaced workers have about 31 fewer employed days (10%) ten years after displacement compared to their matched control workers. Note that employment for displaced as well as control workers decreases after displacement, as by definition both have to be employed in the years before displacement ($t - 4$ to $t - 1$), but thereafter there is no restriction.

Finally, panel D shows that displaced workers receive UI benefits for more days than their matched control workers in the first years after job loss. The difference decreases over time, but a gap remains for all post-displacement years. Note that for displaced as well as for control workers, the days in which they receive UI benefits increase over time, as workers get older. For older workers, entitlement to UI benefit becomes more generous and the likelihood of unemployment increases.

1.4.3 Empirical Specifications

To examine the effect of job loss on retirement, I use quasi-experimental variation in job loss induced by plant closures and compare displaced workers with similar non-displaced workers. Each displaced worker is matched to a non-displaced “twin”, who is of exactly the same age, birth cohort, displacement year and education and who has a similar occupation, wage and pension contribution points before the job loss (see section 1.4.1). This enables me to use the matched “twins” as a control group for the displaced workers and to compare similar workers, with some experiencing the shock and the others not.

Baseline Specification. To show how the overall treatment effect evolves

over age, I run the following specification using the matched sample:

$$y_{i,k} = \sum_{k=55}^{66} \gamma_k \times 1[\text{Age} = k] + \sum_{k=55}^{66} \beta_k \times 1[\text{Age} = k] \times \text{Disp}_i + \epsilon_{i,k}, \quad (1.1)$$

where $y_{i,k}$ is the outcome of worker i at age k (e.g. an indicator equal to 1 if worker i is retired at age k and zero otherwise), 1 is an indicator equal to one at age k and zero otherwise, Disp_i is an indicator equal to 1 if worker i has been displaced and $\epsilon_{i,k}$ is the error term. Standard errors are clustered at the level of matched worker pairs. I use the probability to be retired and estimated pension benefits as main outcome variables.

In this specification, the main coefficients of interest are the β_k 's, which measure the effect of job loss on outcome y (e.g. the probability to be retired) at age k . More precisely, the coefficient β_k estimates the change in y (e.g. the retirement probability) of displaced workers relative to similar non-displaced workers at age k .

In addition to specification (1.1), I estimate the following specification, using larger age bins:

$$y_{i,k} = \sum_{k=1}^3 \gamma_k \times 1[\text{Age period} = k] + \sum_{k=1}^3 \beta_k \times 1[\text{Age period} = k] \times \text{Disp}_i + \epsilon_{i,k}, \quad (1.2)$$

where 1 is an indicator equal to one at age period k and all other variables are defined as in specification (1.1). Age periods are defined as follows: (i) 55-59 and (ii) 60-64 and (iii) 65-66.

Specification for Heterogeneity Analysis. Finally, I investigate how the treatment effect differs by subgroups including age, wage, level of education, economic condition, and occupation. To do so, I estimate the following specification:

$$y_{i,k,g} = \sum_{k=1}^3 \gamma_k \times 1[\text{Age period} = k] \times 1[\text{Group} = g] + \sum_{k=1}^3 \beta_k \times 1[\text{Age period} = k] \times 1[\text{Group} = g] \times \text{Disp}_i + \epsilon_{i,k,g}, \quad (1.3)$$

where $y_{i,k,g}$ is the outcome of worker i at age k belonging to group g ,

$1[\text{Group} = g]$ is an indicator for subgroup g and all other variables are defined as in specification (1.2). Subgroups are measured in pre-displacement year $t - 1$.

Identifying Assumption. For all analyses of specifications (1.1), (1.2) and (1.3), the identifying assumption is that, without the job loss, the change in $y_{i,k}$ would have been comparable between displaced and matched “twin” workers. I cannot test this assumption, but I can show how labor market outcomes of displaced and matched control workers evolved before the job loss. Ideally, the pre-displacement trends of the two groups would be very similar and no significant differences in pre-displacement outcomes would be observed. Figure 1.1 shows that no significant differences between the two groups can be observed before the job, suggesting that the two groups would have followed similar trajectories had the plant closure not taken place²⁰.

1.5 The Effect of Job Loss on Retirement

In this section, I document that displaced workers change their retirement behavior in response to job loss. Despite this change, they still experience significant losses in estimated pension benefits. I first show the overall impact of job loss on retirement timing and estimated pension benefits, and examine how the treatment effect differs for different groups of workers. I then present evidence for the lifetime costs of job loss and show evidence of the robustness of the results. Finally, I compare the results with estimates of previous studies on the impact of late-career job loss.

1.5.1 Overall Effect

Retirement Timing. Panel A of figure 1.2 shows how job loss affects workers’ retirement timing. It shows the results from estimating specification (1.1). The figure plots the difference in the probability of being retired for displaced workers relative to similar non-displaced workers. Panel A shows that the probability of being retired is very similar for displaced and matched control workers before age 60. Coefficients are statistically indistinguishable from

²⁰Note that the plant closure takes place between year $t - 1$ and t , therefore the two groups begin to diverge in year $t - 1$.

zero. After the age of 60, the retirement probability decreases significantly for displaced workers relative to their matched “twin” workers. Displaced workers are about 6 pp (14%) less likely to be retired by the age of 60 compared to similar non-displaced workers. The decrease in the retirement probability of displaced workers is particularly large at age 60-64, when most workers can only retire early. After the age of 64, the difference in the retirement probability of displaced and control workers is small and statistically insignificant²¹. This shows that displaced workers change their retirement behavior in response to job loss by postponing their retirement. Table 1.2 confirms the results using larger age bins (specification (1.2)). Column 1 shows that while job loss does not affect the probability of retirement before age 60, displaced workers are about 6 pp (10%) less likely to be retired between age 60 and 64. Thereafter, the difference in the probability of being retired is close to zero.

Estimated Pension Benefits. Despite the change in retirement timing, displaced workers still experience significant losses in estimated pension benefits. Panel B of figure 1.2 indicates that job loss is associated with significant losses in estimated pension benefits, which are particularly large at age 60/61²². The point estimates are broadly similar at age 63-66, where displaced workers receive approximately 65 € less estimated monthly pension benefits than their matched “twin” workers. Pooling age into larger bins (specification (1.2)), column 2 of table 1.2 shows that job loss leads to about 63 € (5%) and 60 € (5%) lower estimated monthly pension benefits at age 60-64 and age 65-66, respectively. In other words, the average displaced worker who retires at age 65 will have received about 11,000 € lower estimated pension benefits by age 80.

Gains from Adjusting Retirement Behavior. What would estimated pension benefits amount to if displaced workers do not adjust their retirement timing? Answering this question is challenging, as one needs to know the counterfactual (unobserved) retirement behavior of displaced workers if they did not lose their jobs. To overcome this challenge, I take advantage of having

²¹Appendix figure 1.C.6 shows results estimating the effect from the age of job loss until age 66, using an unbalanced panel. The figure shows that before age 60, there are no differences in the retirement probability between the two groups.

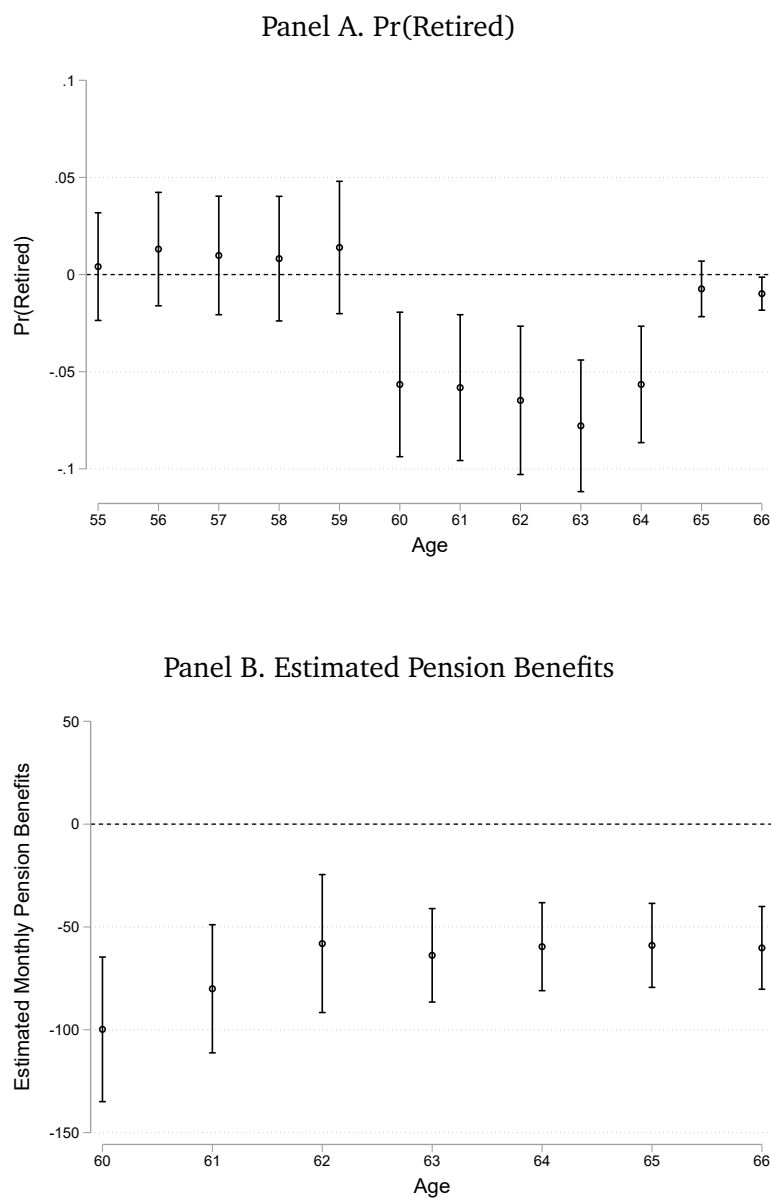
²²Note that the sample consists of workers in which both workers of the matched pair receive pension benefits. As workers are not yet eligible to claim pension benefits before age 60, the treatment effect is estimated from age 60 onwards.

matched each displaced worker with a non-displaced “twin.” I assign each displaced worker the age of retirement entry of the matched “twin” and calculate the estimated pension benefits in this scenario.

Table 1.2: Effect of Job Loss on Retirement - Overall Effect

	(1) Pr(Retired)	(2) Estimated Monthly Pension Benefits	(3) Hypothetical Estimated Monthly Pension Benefits
Age 55-59	0.00983 (0.0147)		
Age 60-64	-0.0628*** (0.0150)	-63.35*** (9.833)	-93.55*** (8.224)
Age 65-66	-0.00861* (0.00520)	-59.58*** (10.20)	-136.6*** (7.518)
Control Mean 55-59	0.188		
Control Mean 60-64	0.598	1172.1	1172.1
Control Mean 65-66	0.982	1193.9	1193.9
# Cluster	1220	1220	1220
Observations	29280	7214	7214

Notes: This table displays the effect of job loss on the retirement probability (1), estimated monthly pension benefit (2) and hypothetical estimated monthly pension benefits (3) of displaced workers relative to matched control workers. Observations are at the worker \times age level. Coefficients are estimated using specification (1.2). Standard errors (in parentheses) are clustered at the level of matched worker pairs. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, respectively.

Figure 1.2: Effect of Job Loss on Retirement - Overall Effect

Notes: This figure shows the effect of job loss on the retirement probability (panel A) and estimated monthly pension benefits (panel B) of displaced workers relative to matched control workers. Observations are at the worker \times age level. Coefficients are estimated using specification (1.1). 95% confidence intervals are derived from standard errors clustered at the level of matched worker pairs.

Column 3 of table 1.2 shows that losses in estimated monthly pension benefits would be much higher had displaced workers not adjusted their retirement behavior. Without adjusting their retirement behavior, displaced workers would receive estimated monthly pension benefits of about 94 € (8%) lower than those of their matched “twin” workers from age 60-64 and about 137 € (11%) lower from age 65/66. Compared to the losses in the case of adjustment, losses in estimated monthly pension benefits of displaced workers would be about 31 € larger from age 60-64 and about 77 € larger from age 65/66. To put this in perspective, the average displaced worker who retires at age 65, can recoup about 25,000 € of estimated pension benefits by age 80 by adjusting his retirement behavior.


1.5.2 Heterogeneity Analysis

The previous section showed that displaced workers react to job loss by delaying retirement. Yet, they still receive significantly lower estimated pension benefits. In this section, I investigate how the treatment effect differs across subgroups²³. For this analysis, I estimate specification (1.3) and show the corresponding coefficient plots.

By Age. The negative consequences of job loss vary considerably by workers’ age at job loss (Couch and Placzek, 2010; Rinz, 2022; Salvanes et al., *ming*). Relative to younger workers, older workers face higher immediate earnings losses due to higher pre-displacement wages and difficulties in finding a new job (Farber et al., 2019; Carlsson and Eriksson, 2019). However, for those workers, reduced earnings occur over a shorter period. Job loss will alter financial incentives to retire through its effect on wages (incentives to continue working) and future pension benefits (incentives to retire). It is not clear how these changes in retirement incentive will differ by workers’ age at job loss and whether the treatment effect is expected to be larger for workers displaced at a young age or later in life.

²³Note that the matching procedure ensures that matched worker pairs are identical with respect to the studied characteristics.

Figure 1.3: Heterogeneous Effects by Age at Job Loss

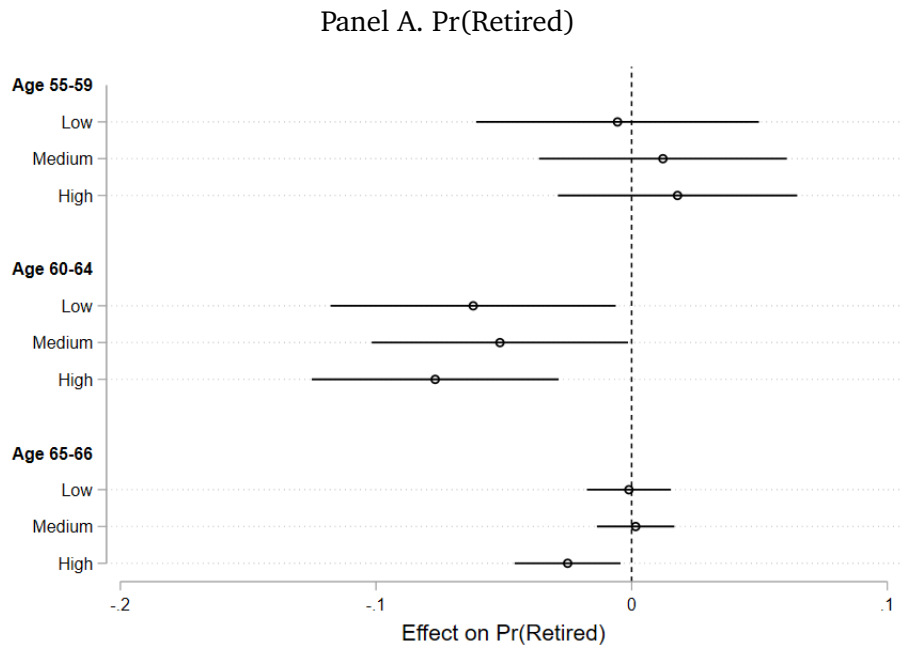


Notes: This figure shows the effect of job loss on the retirement probability (panel A) and estimated monthly pension benefits (panel B) of displaced workers relative to matched control workers by age groups. Age groups are defined in $t - 1$. Observations are at the worker \times age level. Coefficients are estimated using specification (1.3). 95% confidence intervals are derived from standard errors clustered at the level of matched worker pairs.

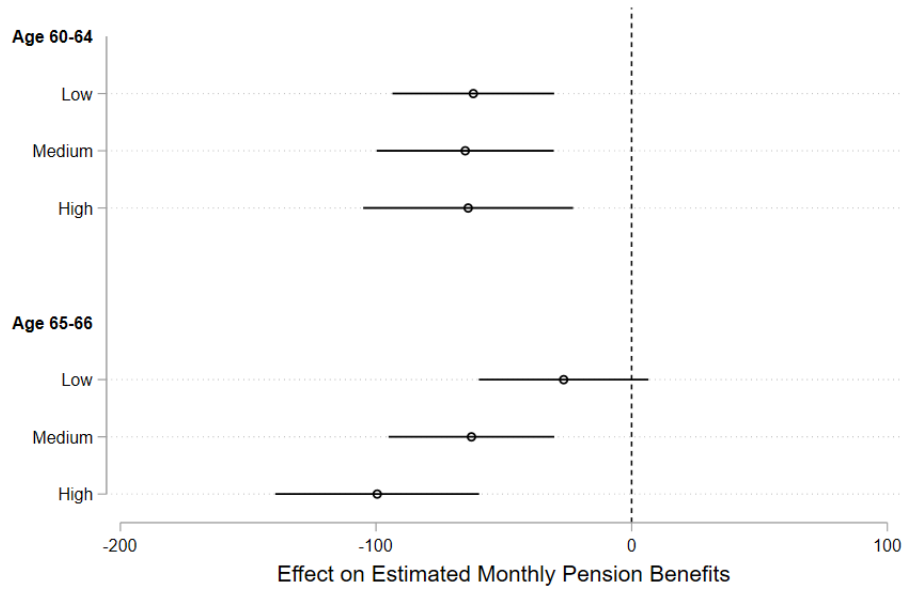
Panel A of figure 2.5 shows that the effect of job loss on the retirement probability varies by workers' age at displacement, although group differences are not statistically significant. In this figure, age is binned into three age groups: (i) 35-40, (ii) 41-45, and (iii) 46-50. Younger workers (35-40) are slightly less likely to be retired compared to their matched control workers from age 55-59, while workers aged 46-50 are more likely to be retired. Some older workers may not be able to find a permanent job after the displacement and therefore decide to permanently exit the labor market even before reaching the statutory retirement age. From age 60 to 64, job loss significantly decreases the retirement probability of young and middle-aged workers. The point estimates imply that job loss decreases the retirement probability of young workers by about 9 pp (15%) and of workers aged 41 to 45 by about 8 pp (14%). For older workers, no significant effect can be observed. After the age of 64, the treatment effect is close to zero for workers of all age groups. Panel B of figure 2.5 shows the impact on estimated pension benefits. Losses in estimated pension benefits are slightly smaller for older displaced workers than for young and middle-aged workers from age 60-64. In contrast, losses are smaller for younger displaced workers from age 65/66. Specifically, while job loss leads to approximately 45 € (4%) lower estimated monthly pension benefits of young workers, losses of middle-aged and older workers are about 62 € (5%) and about 71 € (6%), respectively.

By Pre-displacement Wage. Panel A of figure 1.4 shows how the impact on the retirement probability differs by workers' pre-displacement wage. For this analysis, I consider three wage groups based on tertiles of workers' pre-displacement wage. Panel A shows that job loss is associated with a similar decrease in the retirement probability across wage groups. However, the impact on the estimated pension benefits differs between groups from age 65/66, although differences are not statistically significant (panel B). Displaced workers with low pre-displacement earnings experience losses in estimated monthly pension benefits of approximately 27 € (3%), whereas displaced workers with high pre-displacement earnings lose approximately 100 € (7%).

Figure 1.4: Heterogeneous Effects by Pre-displacement Wage



Panel B. Estimated Pension Benefits

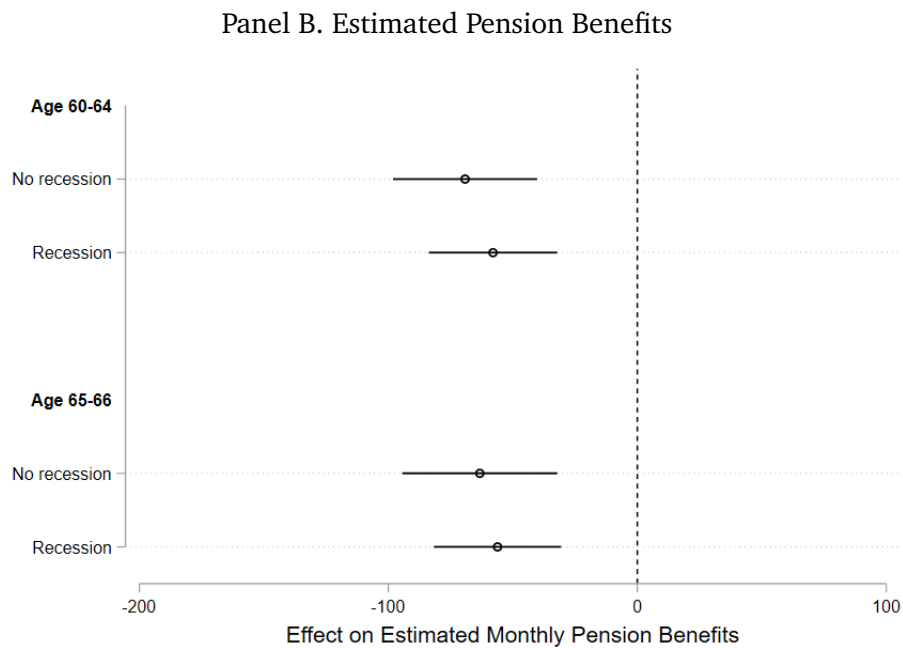



Notes: This figure shows the effect of job loss on the retirement probability (panel A) and estimated monthly pension benefits (panel B) of displaced workers relative to matched control workers by wage groups. Wage groups are defined in $t - 1$. Observations are at the worker \times age level. Coefficients are estimated using specification (1.3). 95% confidence intervals are derived from standard errors clustered at the level of matched worker pairs.

By Economic Conditions at Job Loss. The negative effects of job loss vary by the economic conditions at the time of job loss. Workers who are displaced during recessions experience higher earnings losses than those displaced during normal times (Davis and von Wachter, 2011; Schmieder et al., 2023). In the long term, earnings losses can be explained to a large extent by lower wages, which affect workers' incentives to retire directly and indirectly through its effect on future pensions. To test how the impact on displaced workers' retirement behavior differs by the economic conditions at the time of displacement, I follow Schmieder et al. (2023) and define recessions by changes in the national unemployment rate. Specifically, I consider recessions as years with a year-on-year change in the unemployment rate of -0.5 ²⁴. Panel A of figure 1.5 indicates that the decrease in the retirement probability is larger for workers displaced during recession, from age 55-59 as well as age 60-64. By contrast, there is no evidence of differential impacts on the estimated pension benefits (panel B).

²⁴According to this definition, the following years are classified as recessions: 1981, 1982, 1983, 1991, 1992, 1993, 1994, 1996, 1997, 2002, 2003, 2004, and 2005.

Figure 1.5: Heterogeneous Effects by Economic Conditions at Job Loss




Notes: This figure shows the effect of job loss on the retirement probability (panel A) and estimated monthly pension benefits (panel B) of displaced workers relative to matched control workers by economic conditions at job loss. Economic conditions (recession/no recession) are defined in $t - 1$. Observations are at the worker \times age level. Coefficients are estimated using specification (1.3). 95% confidence intervals are derived from standard errors clustered at the level of matched worker pairs.

By Education. Panel A of figure 1.6 shows that the impact of job loss on the retirement probability differs by workers' education level²⁵. The decline in the retirement probability relative to similar non-displaced workers appears to be larger for workers with no/unrecognized education from age 60-64. Note that these estimates are quite imprecise, as about 90% of workers in the analysis sample have completed vocational training or a high school diploma, while the remaining 10% either have no education or a university degree. Panel B shows that the negative effect on estimated pension benefits is larger for workers with a university degree, but again estimates are quiet imprecise.

²⁵I consider the following three education groups: (i) no/unrecognized education, (ii) vocational training/high school diploma, and (iii) university degree, which I measure in $t - 1$.

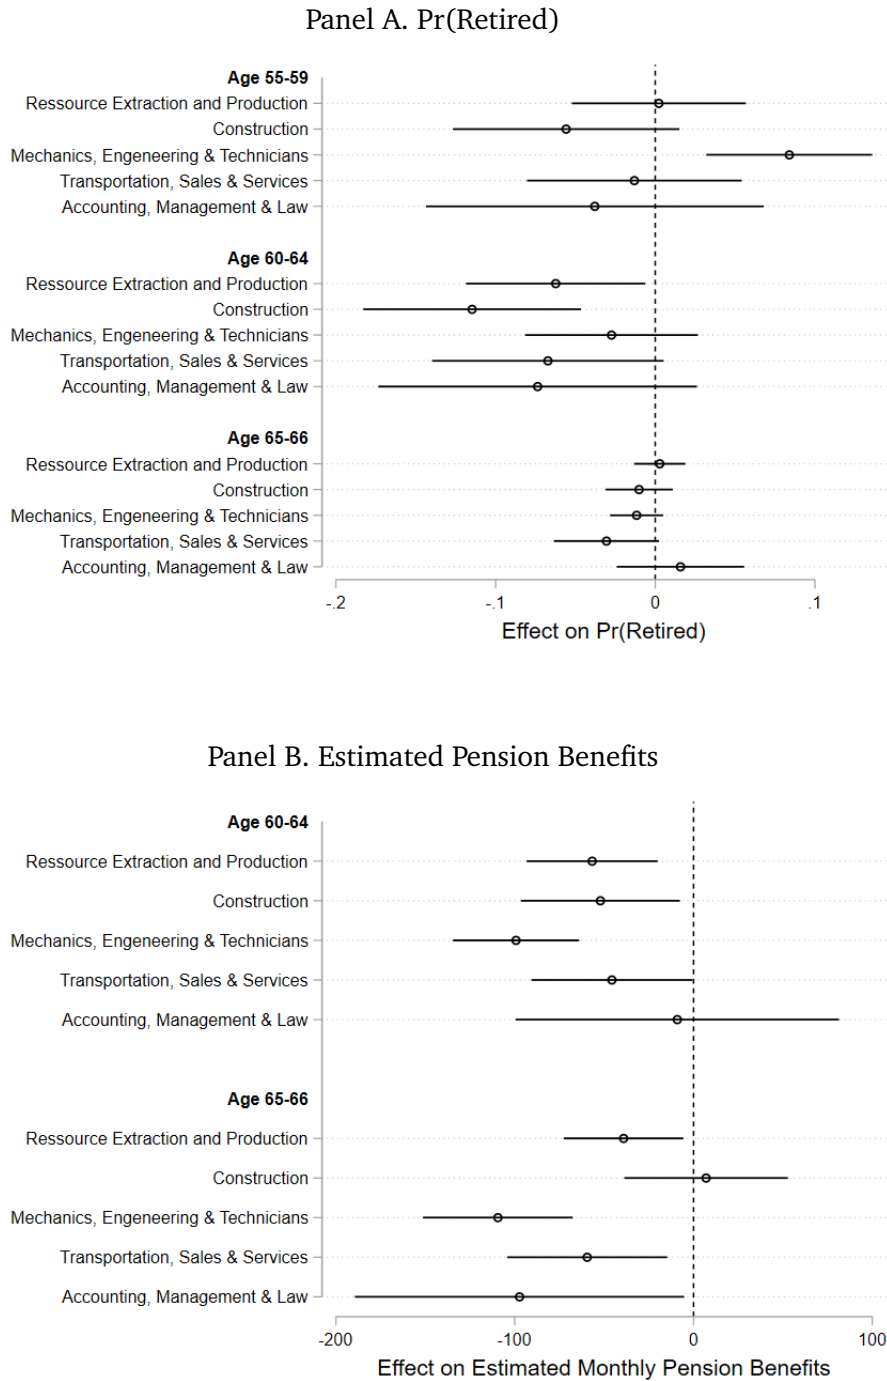
Figure 1.6: Heterogeneous Effects by Pre-displacement Education



Notes: This figure shows the effect of job loss on the retirement probability (panel A) and estimated monthly pension benefits (panel B) of displaced workers relative to matched control workers by education. Education groups are defined in $t - 1$. Observations are at the worker \times age level. Coefficients are estimated using specification (1.3). 95% confidence intervals are derived from standard errors clustered at the level of matched worker pairs.

By Pre-displacement Occupation. The treatment effect may also differ by workers' occupation. Panel A of figure 1.7 shows that the decrease in the retirement probability of displaced workers relative to non-displaced workers differs across occupations and is particularly large for worker in construction. In contrast, losses in estimated pension benefits (panel B) are particularly large for workers in high-paying occupations. For example, losses are particularly large for workers in mechanics and engineering, and accounting, management and law from age 60-64.

Figure 1.7: Heterogeneous Effects by Pre-displacement Occupation



Notes: This figure shows the effect of job loss on the retirement probability (panel A) and estimated monthly pension benefits (panel B) of displaced workers relative to matched control workers by occupation groups. Occupation groups are defined in $t - 1$. Observations are at the worker \times age level. Coefficients are estimated using specification (1.3). 95% confidence intervals are derived from standard errors clustered at the level of matched worker pairs.

1.5.3 The Lifetime Costs of Job Loss

Displaced workers experience large losses during their working lives, but also beyond. To provide summary measures of the *lifetime costs* of job loss, I compute the PDV of lifetime losses for displaced workers relative to their matched control workers for different outcomes. To do so, I start by computing the PDV for each worker i :

$$PDV(y_i) = \sum_{t=D}^T \frac{1}{(1+r)^{(t-D)}} y_i,$$

where y is one of the following outcome variables: labor earnings, estimated pension benefits, or income. Income includes labor earnings, UI benefits, and estimated pension benefits. D is the year of displacement, T the year of death, and r the discount rate. As workers can only be observed until the age of 66, I make the assumption that all workers will receive pension benefits until the age of 80. I assume a discount rate of 3%.

I then simply use the values of $PDV(y_i)$ and compute the average difference between displaced ($d = 1$) and matched non-displaced workers ($d = 0$):

$$\Delta PDV(y) = PDV(y_1) - PDV(y_0).$$

Overall Effect. Panel A of table 1.3 shows the results for the overall sample. Column 1 shows that job loss leads to sizeable losses in the PDV of income. On average, job loss leads to losses in the PDV of income of about 127,000 (€16%). Column 2 shows that the PDV of earnings of displaced workers is about 125,000 (€19%) lower compared to non-displaced workers. Relative to the losses in the PDV of income, the percentage loss in the PDV of earnings is larger, because income includes earnings as well as UI benefits and estimated pension benefits. Column 3 shows that job loss reduces the PDV of estimated pension benefits by about 8,000 (€8%).

Table 1.3: Life-time Costs of Job Loss

	(1) PDV Income (incl. Estimated Pension Benefits)	(2) PDV Labor Earnings	(3) PDV Estimated Pension Benefits
<i>Panel A. Overall</i>			
	-127161.7*** (11632.9)	-125282.0*** (11043.7)	-8434.8*** (857.4)
Control Mean	780690.6	664929.7	110294.5
% of Mean	-16.29	-18.84	-7.65
Observations	1220	1220	1220
<i>Panel B. Age 30-39</i>			
	-101138.4*** (23526.7)	-99385.2*** (22197.2)	-5878.0*** (1614.9)
Control Mean	852025.0	757524.6	89267.3
% of Mean	-11.87	-13.12	-6.66
Observations	395	395	395
<i>Panel C. Age 40-45</i>			
	-147289.6*** (20418.0)	-144857.8*** (19530.6)	-8928.9*** (1496.1)
Control Mean	821828.6	705314.7	110499.4
% of Mean	-17.92	-20.54	-8.08
Observations	373	373	373
<i>Panel D. Age 45-50</i>			
	-133293.2*** (16696.5)	-131758.7*** (15870.6)	-10261.3*** (1351.9)
Control Mean	684403.8	550684.9	128500.9
% of Mean	-19.48	-23.93	-7.99
Observations	452	452	452
<i>Panel E. Hypothetical</i>			
	-178377.3*** (15618.0)	-171902.5*** (14903.1)	-11848.4*** (1380.2)
Control Mean	780690.6	664929.7	110294.5
% of Mean	-22.8	-25.9	-10.7
Observations	1220	1220	1220

Notes: This table shows the present discounted value (PDV) of income, labor earnings and estimated pension benefits for displaced workers relative to similar non-displaced workers. Income includes labor earnings, UI benefits and estimated pension benefits. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, respectively.

By Age at Job Loss. Panel B, C, and D of table 1.3 show that the negative impact of job loss is large across all age groups, but the percentage loss in the PDV differs. Workers aged 35 to 40 experience the smallest loss in the PDV of income and earnings. For these workers, the PDV of income is about 12% lower relative to similar non-displaced workers (panel B, column 1), while it is about 18% lower for workers aged 41 to 45 (panel C, column 1) and 19% lower for workers aged 46 to 50 (panel D, column 1). In terms of the PDV of estimated pension benefits, losses are again the smallest in the 35-40 age bracket (panel B, column 3) and similar for workers aged 41 to 45 (panel C, column 3) and 46 to 50 (panel D, column 3). The youngest age group experiences losses in the PDV of estimated pension benefits of about 8%, while it is about 9% for the other groups.

By Subgroups. I also examine how the lifetime costs of job loss differ by further subgroups. Appendix table 1.C.6 summarizes the results separately in subgroups for education, pre-displacement wage, pre-displacement occupation, and economic conditions at the time of job loss²⁶. The results show that the lifetime costs of job loss are particularly large for workers with high pre-displacement wages (panel C) and those with a university degree (panel F). While workers with high pre-displacement wages experience losses in the PDV of income of about 21%, losses are only about 11% for workers with low pre-displacement earnings. Losses also vary by occupation, with the largest losses occurring among workers in high-paying occupations. Workers in mechanics and engineering (panel K) and accounting, management and law (panel M) experience losses in the PDV of income of about 22% and 21%, while construction workers experience losses of only about 6%.

Gains from Adjusting Retirement Behavior. If displaced workers do not adjust their retirement timing, what would their lifetime losses amount to? To answer this question, I follow the same strategy as in section 1.5.1 and simply assign each displaced worker the age of retirement entry of his matched “twin” worker and then calculate the PDV in this scenario. Panel E of table 1.3 shows that lifetime losses would be much bigger if displaced workers did not adjust their retirement timing. Displaced workers would have losses in the PDV of income of about 178,000 € (23%) (panel E, column 1), compared to

²⁶Subgroups are defined as in section 1.5.2.

about 127,000 € (16%) in the case of adjustment (panel A, column 1). Put differently, by adjusting their retirement timing displaced workers can recoup about 51,000 € (7%) in terms of the PDV of income. Overall, the results highlight that the costs of job loss are enormous and more far-reaching than short-term or medium-term estimates of earnings losses can show.

1.5.4 Robustness Analysis

I examine the robustness of the results to alternative samples, a different matching specification and an alternative estimation strategy. The results are displayed in table 1.4. Panel A shows the results for the impact of job loss on the probability of being retired and panel B on the estimated pension benefits. For comparison, column 1 shows the baseline results.

Table 1.4: Effect of Job Loss on Retirement - Robustness Checks

	(1) Baseline	(2) Alternative Sample	(3) Without exits before 55	(4) Other Matching	(5) Pre-displacement Controls
<i>Panel A. Pr(Retired)</i>					
Age 55-59	0.00984 (0.0147)	0.00968* (0.00552)	0.0153* (0.00921)	-0.00422 (0.00830)	0.00404 (0.0104)
Age 60-64	-0.0628*** (0.0150)	-0.0516** (0.0215)	-0.0670*** (0.0156)	-0.0654*** (0.0154)	-0.0616*** (0.0108)
Age 65-66	-0.00861* (0.00520)	0.00538 (0.0109)	-0.0166*** (0.00724)	0.0180 (0.0146)	0.00526 (0.00490)
Control Mean 55-59	0.188	0.00430	0.0666	0.0626	0.206
Control Mean 60-64	0.598	0.365	0.529	0.460	0.611
Control Mean 65-66	0.982	0.965	0.972	0.837	0.981
# Cluster	1220	372	1057	1058	12471
Observations	29280	8928	25368	27312	149652
<i>Panel B. Estimated Monthly Pension Benefits</i>					
Age 60-64	-63.35*** (9.833)	-56.44*** (20.42)	-64.53*** (10.08)	-48.71*** (13.67)	-63.66*** (6.714)
Age 65-66	-59.58*** (10.20)	-59.02*** (15.34)	-70.70*** (10.10)	-33.16*** (11.38)	-61.63*** (7.515)
Control Mean 60-64	1172.1	1306.3	1195.3	1172.2	1189.2
Control Mean 65-66	1193.9	1353.0	1254.8	1164.4	1218.1
# Cluster	1220	372	1057	1058	12471
Observations	7214	2084	6540	5856	48272

Notes: This table shows the effect of job loss on the retirement probability (panel A) and estimated monthly pension benefits (panel B) of displaced workers relative to matched control workers. Observations are at the worker \times age level. Coefficients are estimated using specification (1.2). 95% confidence intervals are derived from standard errors clustered at the level of matched worker pairs (columns (1)-(4)) or at the worker level (column (5)). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, respectively.

Alternative Samples. A caveat of the data is that it does not include information on pension claims and benefits. The year of retirement entry is defined as the year in which a worker is last observed in the data as either employed or unemployed. In addition to retirement, there are other reasons why people may leave the data, such as becoming self-employed, taking a civil service, or emigrating. For these workers, the year of retirement entry would not be correctly identified. As a robustness check, column 2 presents the results for an alternative sample that only includes workers who exit the data due to retirement. While the data does not provide information on pension claims, information on the reason for the termination of the employment or unemployment spell is included. The alternative sample contains workers for whom the reason for terminating the employment or unemployment period is given as retirement²⁷. In principle, workers in this sample can still permanently exit employment in the year of displacement and use UI benefits as a source of income until they decide to retire. Because the sample size is much smaller than the main sample and workers in this sample are observed until retirement, I only use this sample as a robustness check. Column 2 of table 1.4 shows that the results for this sample are comparable to those of the baseline sample (column 1). The impact on estimated retirement benefits is also similar (panel B, column 2).

As a further robustness check, I consider restricting the sample to workers who do not exit the data before age 55 (column 3). For this sample, the probability of being retired before age 55 is zero. Column 3 of table 1.4 shows that the results for this sample are quite similar to those of the baseline sample.

Alternative Matching Specification. Column 4 of table 1.4 shows the results using an alternative matching specification. Specifically, I change the set of covariates used in the matching procedure and match on the following covariates: displacement year (1978-2004), birth cohort (1939-1954), age (35-50), industry (8 groups)²⁸, tenure (2 groups) (all measured in $t - 1$) and

²⁷Note that most employers do not further specify the reason for the end of the employment relationship, but rather report “Deregistration due to end of employment.” Since many employers do not report the retirement entry of their employees, the sample size is much smaller than the main sample.

²⁸I consider the following industries: (i) agriculture and mining, (ii) manufacturing, (iii) electricity, water and waste management, (iv) construction, (v) trade and transport, (vi)

log wages (2 groups) (measured in $t - 3$ and $t - 4$). The results are similar to the baseline matching specification, with slightly smaller losses in estimated pension benefits (panel B, column 4).

Pre-displacement Controls. In the main analysis, I apply matching to find appropriate counterfactuals for the displaced workers. Each displaced worker is matched to a non-displaced “twin.” In column 5 of table 1.4, I show that the results are robust to using an alternative estimation strategy. In particular, I use all non-displaced workers as a control group and control for pre-displacement characteristics rather than using the matched control group. The coefficients in column 5 are obtained from estimating the following specification:

$$y_{i,k} = \sum_{k=1}^3 \gamma_k \times 1[\text{Age period} = k] + \sum_{k=1}^3 \beta_k \times 1[\text{Age period} = k] \times \text{Disp}_i + X_i + \epsilon_{i,k},$$

where $y_{i,k}$ is the outcome of worker i at age k , Disp_i is an indicator equal to 1 for displaced workers and zero otherwise, X_i are pre-displacement worker characteristics and age periods are defined as in specification (1.2). Standard errors are clustered at the worker level.

I use the same set of covariates as in the matching procedure as control variables (X_i). All control variables are used as continuous variables, if possible, and measured at $t - 1$. Note that the sample size is larger than the baseline sample, since all non-displaced workers are used as a control group rather than using only one matched control worker. Column 6 shows that the results from this alternative estimation strategy are quite similar to those using the matched sample.

1.5.5 Comparison with Effect of Late-Career Job Loss

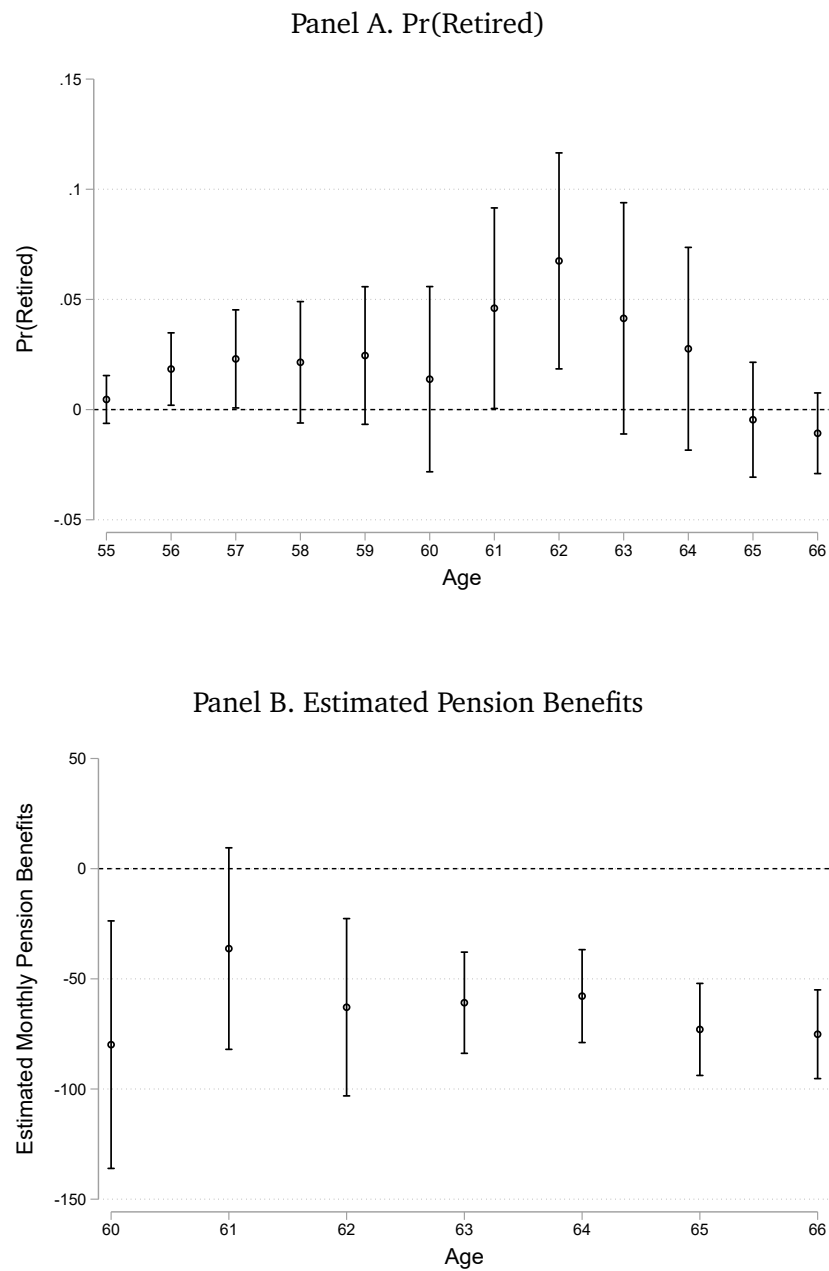
Some previous papers have examined how late-career job loss affects retirement (Chan and Huff Stevens, 2001; Chan and Stevens, 2004; Merkurieva, 2019). They show that late-career job loss substantially increases the rate of exit from employment (Chan and Huff Stevens, 2001) and results in earlier retirement (Chan and Stevens, 2004; Merkurieva, 2019). For example, Chan

services, (vii) banking, insurance and real estate, (viii) science, education and health

and Stevens (2004) estimate that a late-career job loss doubles the probability of retirement each year. This may be because older displaced workers face higher costs of job search, find it more difficult to find a new job, or have the opportunity to use UI benefits as a stepping stone into early retirement.

For comparison with previous studies, figure 1.8 shows how late-career job loss affects retirement. Specifically, it shows the results from estimating specification (1.1) for a sample of workers with a late-career job loss²⁹. A late-career job loss is defined as a job loss that occurs at age 50 or later (Chan and Huff Stevens, 2001). Panel A of figure 1.8 shows that following a late-career job loss, displaced workers are more likely to be retired than their non-displaced “twin” workers before age 65. The probability of being retired is about 2 pp (5%) higher at ages 55-59 and about 4 pp (9%) higher at ages 60-64, which is in line with previous studies. Thereafter, the estimated treatment effect is close to zero. In contrast to Chan and Stevens (2004), I estimate that late-career job loss also affects estimated pension benefits. Panel B shows that displaced workers receive about 74 € (6%) lower estimated pension benefits. This may be explained by differences in pension systems across countries. In contrast to several other countries, all contribution years are taken into account in Germany when calculating pension benefits. Compared to systems where benefits are calculated based on a worker’s best years or last years before retirement, the losses in pension benefits are larger.

²⁹Appendix table 1.C.7 shows results for larger age bins (specification (1.2)).

Figure 1.8: Effect of Job Loss on Retirement - Late Career Job Loss

Notes: This figure shows the effect of job loss on the retirement probability (panel A) and estimated monthly pension benefits (panel B) of displaced workers relative to matched control workers. Observations are at the worker \times age level. Coefficients are estimated using specification (1.1). 95% confidence intervals are derived from standard errors clustered at the level of matched worker pairs.

1.6 Potential Channels


A central result of the paper is that job loss leads displaced workers to delay retirement. Displaced workers are significantly less likely to be retired from age 60-64. What mechanisms can explain this finding? In this section, I present evidence for different channels.

1.6.1 Pathways into Retirement

Displaced workers can delay retirement by working longer or by taking up other social insurance programs, such as UI benefits – often referred to as *program substitution* (Inderbitzin et al., 2016). In both cases, future pension benefits will increase and displaced workers will avoid the penalty for early pension claiming if delaying retirement entry until the NRA. Future pension benefits will increase more if they stay in employment longer, since pension contributions from wage income are higher than those from UI benefits.

To show how much of the decrease in the retirement probability is driven by increases in employment and how much due to taking alternative routes to retirement, figure 1.9 shows the impact of job loss on the probability of being retired, along with the probability of being employed, unemployed, or inactive (temporary out of the labor force).³⁰ More specifically, it shows the results when specification (1.1) is estimated separately for the different outcomes.

³⁰Outcome variables are defined as in section 1.3.2.

Figure 1.9: Pathways into Retirement

Notes: This figure shows the effect of job loss on the probability to be retired, employed, unemployed or inactive of displaced workers relative to matched control workers. Observations are at the worker \times age level. Coefficients are estimated using specification (1). 95% confidence intervals are derived from standard errors clustered at the level of matched worker pairs.

Figure 1.9 shows that the probability of being retired is close to zero before age 60. However, job loss has a significant impact on the probability of being employed, unemployed or inactive. Displaced workers are less likely to be employed and more likely to be unemployed or inactive compared to their non-displaced “twins” before the age of 60. Note that the negative effect on the employment probability is decreasing from age 55-59, driven by decreases in the employment probability of non-displaced workers (see appendix figure 1.C.4). After the age of 60, displaced workers delay retirement and the retirement probability decreases significantly relative to similar non-displaced workers. This decrease can be explained by an increase in employment as well as unemployment, with the effect strongly depending on workers’ age at job loss.

Appendix figure 1.C.5 shows the results for the different age groups. The decrease in retirement probability from age 60 to 64 can mainly be explained by increases in employment for younger workers (panel A) and a higher probability of being unemployed for older workers (panel C). Specifically, while the decrease in the retirement probability (from age 60-64) of younger workers can be explained by an increase in employment of about 6 pp (19%), it amounts to an increase of only 3 pp (10%) for workers aged 41 to 45 and a zero effect for workers aged 46 to 50. Overall, younger workers tend to delay retirement by working longer, while older workers use UI benefits as a stepping stone into retirement.


1.6.2 Alternative Channels

Effect on Health. Displacement may also influence workers’ retirement decision via changes in their health status, e.g. through changes in workers’ life expectancy or their ability to work in old age. In general, evidence for the effects of job loss on health is mixed. While some studies find only limited evidence of a negative impact of job loss on health (Black et al., 2015; Browning and Heinesen, 2012; Schaller and Stevens, 2015), others show large adverse effects, with a sharp increase in mortality after displacement (Sullivan and Von Wachter, 2009).

To investigate the health effects of job loss, I use the number of sick leave days as a proxy for workers’ health status, similar to Staubli and Zweimüller

(2013). Since the data is based on employment records submitted by employers as well as records from employment agencies, they unfortunately do not include direct information on workers' health status. The number of sick leave days is therefore the only possible proxy. Sick leave days can be identified, as employers and employment agencies report sick leave days of more than six weeks or hospital stays (of any length) to the social security system. Note that the observed outcome is conditional on workers participating in the labor market, either working or receiving UI benefits. Figure 1.10 displays the effect of job loss on the yearly number of sick leave days, using specification (1.1). It shows only a very small increase in the number of sick leave days of displaced workers relative to similar non-displaced workers. The effect is however very small.

Figure 1.10: Effect of Job Loss on Number of Sick Leave Days



Notes: This figure shows the effect of job loss on the number of sick leave days of displaced workers relative to matched control workers. Observations are at the worker \times age level (from age 55-66). Coefficients are estimated using specification (1.1). 95% confidence intervals are derived from standard errors clustered at the level of matched worker pairs.

1.7 Conclusion

Economists have long been interested in studying the effect of job loss. The literature has shown that job loss leads to large and persistent earnings losses (Couch and Placzek, 2010; Davis and von Wachter, 2011; Jarosch, 2023; Jacobson et al., 1993; Lachowska et al., 2020; Schmieder et al., 2023) and impacts workers on various other dimensions as well (Del Bono et al., 2012; Huttunen et al., 2018, 2011; Khanna et al., 2021; Sullivan and Von Wachter, 2009). Yet, there is little evidence of the impact of job loss on retirement. Some studies have examined how late-career job loss affects workers' retirement behavior (Chan and Huff Stevens, 2001; Chan and Stevens, 2004; Merkurieva, 2019), but there is no evidence of the impact of job loss on younger and middle-aged workers.

Through this study, I provide first evidence of the long-term effects of job loss on retirement. Exploiting quasi-random variation in job losses induced by plant closures, I document three new findings. First, I show that displaced workers react to the shock by delaying retirement. At age 60 to 64, younger and middle-aged displaced workers are about 10% less likely to be retired relative to similar non-displaced workers. This effect differs across subgroups and is larger for younger and low-skilled workers, and for workers displaced during a recession. Even if displaced workers delay retirement, they experience losses in their estimated monthly pension benefits. Second, on average, job loss results in about 6% lower estimated monthly pension benefits.

Third, I show that the lifetime costs of job loss are substantial. Displaced workers experience losses in the PDV of income of about 16%. Using non-displaced workers as counterfactuals, I show that losses would have been much bigger had displaced workers not adjusted their retirement behavior. By delaying retirement, an average displaced worker can recoup about 7% in terms of the PDV of income.

Bibliography

- Black, S. E., Devereux, P. J., and Salvanes, K. G. (2015). Losing heart? the effect of job displacement on health. *ILR Review*, 68(4):833–861.
- Börsch-Supan, A. H. and Wilke, C. B. (2004). The german public pension system: how it was, how it will be.
- Brown, K. M. (2013). The link between pensions and retirement timing: Lessons from california teachers. *Journal of Public Economics*, 98:1–14.
- Browning, M. and Heinesen, E. (2012). Effect of job loss due to plant closure on mortality and hospitalization. *Journal of health economics*, 31(4):599–616.
- Card, D., Heining, J., and Kline, P. (2013). Workplace Heterogeneity and the Rise of West German Wage Inequality*. *The Quarterly Journal of Economics*, 128(3):967–1015.
- Card, D., Maestas, N., and Purcell, P. J. (2014). Labor market shocks and early social security benefit claiming. *Michigan Retirement Research Center Research Paper No. WP*, 317.
- Carlsson, M. and Eriksson, S. (2019). Age discrimination in hiring decisions: Evidence from a field experiment in the labor market. *Labour Economics*, 59:173–183. Special Issue on “European Association of Labour Economists, 30th annual conference, Lyon, France, 13-15 September 2018.
- Chan, S. and Huff Stevens, A. (2001). Job loss and employment patterns of older workers. *Journal of Labor Economics*, 19(2):484–521.
- Chan, S. and Stevens, A. H. (2004). How does job loss affect the timing of retirement? *Contributions in Economic Analysis Policy*, 3(1):1–24.
- Coile, C. C. and Levine, P. B. (2007). Labor market shocks and retirement: Do government programs matter? *Journal of Public Economics*, 91(10):1902–1919.

- Couch, K. A. and Placzek, D. W. (2010). Earnings losses of displaced workers revisited. *American Economic Review*, 100(1):572–89.
- Dauth, W. and Eppelsheimer, J. (2020). Preparing the sample of integrated labour market biographies (siab) for scientific analysis: a guide. *Journal for Labour Market Research*, 54(10).
- Davis, S. J. and von Wachter, T. (2011). Recessions and the costs of job loss. *Brookings Papers on Economic Activity*, pages 1–73.
- Del Bono, E., Weber, A., and Winter-Ebmer, R. (2012). Clash of career and family: Fertility decisions after job displacement. *Journal of the European Economic Association*, 10(4):659–683.
- Deutsche Rentenversicherung (2014). Minijob–midijob: Bausteine für die rente.
- Dustmann, C., Ludsteck, J., and Schönberg, U. (2009). Revisiting the German Wage Structure*. *The Quarterly Journal of Economics*, 124(2):843–881.
- Farber, H. S. (2017). Employment, hours, and earnings consequences of job loss: Us evidence from the displaced workers survey. *Journal of Labor Economics*, 35(S1):S235–S272.
- Farber, H. S., Herbst, C. M., Silverman, D., and von Wachter, T. (2019). Whom do employers want? the role of recent employment and unemployment status and age. *Journal of Labor Economics*, 37(2):323–349.
- Geyer, J. and Welteke, C. (2021). Closing routes to retirement for women how do they respond? *Journal of Human Resources*, 56(1):311–341.
- Giesecke, M. and Kind, M. (2013). Bridge unemployment in germany: Response in labour supply to an increased early retirement age. *Ruhr Economic Paper*, (410).
- Gudgeon, M., Schmieder, J. F., Trenkle, S., and Ye, H. (2019). The labor supply effects of unemployment insurance for older workers. Technical report, Mimeo.

- Gudgeon, M. and Trenkle, S. (forthcoming). The speed of earnings responses to taxation and the role of firm labor demand. *Journal of Labor Economics*.
- Gustman, A. L., Steinmeier, T. L., et al. (1986). A structural retirement model. *Econometrica*, 54(3):555–584.
- Hairault, J.-O., Sopraseuth, T., and Langot, F. (2010). Distance to Retirement and Older Workers' Employment: The Case for Delaying the Retirement Age. *Journal of the European Economic Association*, 8(5):1034–1076.
- Hethey-Maier, T. and Schmieder, J. F. (2013). Does the Use of Worker Flows Improve the Analysis of Establishment Turnover? Evidence from German Administrative Data. *Schmollers Jahrbuch : Journal of Applied Social Science Studies / Zeitschrift für Wirtschafts- und Sozialwissenschaften*, 133(4):477–510.
- Huttunen, K., Møen, J., and Salvanes, K. G. (2011). How destructive is creative destruction? effects of job loss on job mobility, withdrawal and income. *Journal of the European Economic Association*, 9(5):840–870.
- Huttunen, K., Møen, J., and Salvanes, K. G. (2018). Job loss and regional mobility. *Journal of Labor Economics*, 36(2):479–509.
- Iacus, S. M., King, G., and Porro, G. (2011). Multivariate matching methods that are monotonic imbalance bounding. *Journal of the American Statistical Association*, 106(493):345–361.
- Iacus, S. M., King, G., and Porro, G. (2012). Causal inference without balance checking: Coarsened exact matching. *Political analysis*, 20(1):1–24.
- Inderbitzin, L., Staubli, S., and Zweimüller, J. (2016). Extended unemployment benefits and early retirement: Program complementarity and program substitution. *American Economic Journal: Economic Policy*, 8(1):253–88.
- Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993). Earnings losses of displaced workers. *The American economic review*, pages 685–709.
- Jarosch, G. (2023). Searching for job security and the consequences of job loss. *Econometrica*, 91(3):903–942.

- Khanna, G., Medina, C., Nyshadham, A., Posso, C., and Tamayo, J. (2021). Job loss, credit, and crime in Colombia. *American Economic Review: Insights*, 3(1):97–114.
- Lachowska, M., Mas, A., and Woodbury, S. A. (2020). Sources of displaced workers' long-term earnings losses. *American Economic Review*, 110(10):3231–66.
- Manoli, D. and Weber, A. (2016). Nonparametric evidence on the effects of financial incentives on retirement decisions. *American Economic Journal: Economic Policy*, 8(4):160–182.
- Merkurieva, I. (2019). Late career job loss and the decision to retire. *International Economic Review*, 60(1):259–282.
- Mika, T. and Krickl, T. (2020). Entwicklung des Übergangs in die Altersrente bei den Geburtsjahrgängen 1936 bis 1952.
- Pfann, G. A. and Hamermesh, D. S. (2001). Two-sided learning, labor turnover and displacement. Working Paper 8273, National Bureau of Economic Research.
- Rentenversicherung, D. (2020). Die richtige Altersrente für Sie.
- Rentenversicherung, D. (2021). Zahlen und Tabellen der gesetzlichen Rentenversicherung – Werte West (ohne Knappschafft) - 1.1. – 30.6.2021. Technical report.
- Rinz, K. (2022). Did Timing Matter? Life Cycle Differences in Effects of Exposure to the Great Recession. *Journal of Labor Economics*, 40(3):703–735.
- Rust, J. and Phelan, C. (1997). How social security and Medicare affect retirement behavior in a world of incomplete markets. *Econometrica*, 65(4):781–831.
- Salvanes, K. G., Willage, B., and Willén, A. L. (forthcoming). The effect of labor market shocks across the life cycle. *Journal of Labor Economics*.

- Schaller, J. and Stevens, A. H. (2015). Short-run effects of job loss on health conditions, health insurance, and health care utilization. *Journal of health economics*, 43:190–203.
- Schmieder, J. F., Von Wachter, T., and Heining, J. (2023). The costs of job displacement over the business cycle and its sources: evidence from germany. *American Economic Review*, 113(5):1208–1254.
- Schwerdt, G. (2011). Labor turnover before plant closure: “leaving the sinking ship” vs. “captain throwing ballast overboard”. *Labour Economics*, 18(1):93–101.
- Seibold, A. (2021). Reference points for retirement behavior: Evidence from german pension discontinuities. *American Economic Review*, 111(4):1126–65.
- Staubli, S. and Zweimüller, J. (2013). Does raising the early retirement age increase employment of older workers? *Journal of public economics*, 108:17–32.
- Stevens, A. H. (1997). Persistent effects of job displacement: The importance of multiple job losses. *Journal of Labor Economics*, 15(1):165–188.
- Stevens, A. H. (2001). Changes in earnings instability and job loss. *ILR Review*, 55(1):60–78.
- Stock, J. H. and Wise, D. A. (1990). Pensions, the option value of work, and retirement. *Econometrica*, 58(5):1151–1180.
- Sullivan, D. and Von Wachter, T. (2009). Job displacement and mortality: an analysis using administrative data. *The Quarterly Journal of Economics*, 124(3):1265–1306.
- US Bureau of Labor Statistics (2022). Worker displacement: 2019-2021.
- Ye, H. (2022). The effect of pension subsidies on the retirement timing of older women. *Journal of the European Economic Association*, 20(3):1048–1094.

Appendix

1.A Institutional Setting

This section provides additional information on the institutional setting in Germany. For more details on the public pension system, see Börsch-Supan and Wilke (2004).

1.A.1 Pension Benefit Calculation

In Germany, monthly pension benefits are calculated according to the following formula:

$$B_i(R_i) = \sum_{t=0}^{R_i-1} \frac{w_{it}}{\bar{w}_t} \times a(R_i) \times PV,$$

where $B_i(R_i)$ are the monthly pension benefits of worker i retiring in year R_i . w_{it} is worker i 's individual wage in year t , \bar{w}_t the average wage of all insured in year t , $a(R_i)$ the adjustment factor and PV the pension value.

This formula has three components: the first component is the *sum of contribution points*, which is the main determinant of a worker's pension benefits. Essentially, for each year of contribution, a worker accumulates some contribution points, which are determined by the individual income w_{it} relative to the average income of all insured (\bar{w}_t). For example, a worker will receive exactly one contribution point in a year if their average yearly income equals the average yearly income of all insured individuals. The sum of contribution points is then the total number of contribution points a worker earned during their working life. Note that pension contributions cannot only be acquired during periods of employment, but also during other insurance periods such as unemployment or illness. However, pension contribution points will be lower during those periods, as income will be lower than income from employment.

The second part of the formula is the *adjustment factor* ($a(R_i)$) which depends on a worker's age when claiming pension. Workers have to bear a penalty for claiming early, whereas they receive a reward if they retire after the NRA. Specifically, a deduction of 0.3% is imposed for each month retiring

before the NRA and a reward of 0.5% is paid for each month retiring after the NRA. For a worker who retires exactly at the NRA, the adjustment factor equals one.

The third part of the formula is the *current pension value*. It translates the adjusted contribution points into monthly pension benefits. In 2021, one contribution point amounted to 34.19f€ or West Germany and 33.47f€ or East Germany (Rentenversicherung, 2021).

1.A.2 Pathways

Over the observation period, the German public pension system has different pathways through which workers can enter retirement, which vary in the *statutory retirement ages* and *requirements*. Table 1.A.1 summarizes the requirements for the different pathways. To claim a regular pension, workers only need a minimum of five contribution years, but there is no possibility of early retirement in this pathway. For all other pathways, the eligibility requirements are stricter than for claiming a regular pension, but each pathway offers an opportunity for early retirement. In case of early retirement, a worker has to bear a penalty of 0.3% for each month retiring before the NRA. Table 1.A.2 shows the statutory retirement ages corresponding to the different pathways. The statutory retirement ages vary by cohort due to pension reforms. In table 1.A.2, the NRA is shown along with the ERA in brackets (if the respective pathway offers an early retirement option). Note that while the ERA in the unemployment pathway was increased by cohort, a special rule, “*Vertrauensschutz*”, meaning “protection of trust”, allowed all workers born before 1952 to still go into early retirement at age 60 if they fulfilled one of the two criteria : (i) they were registered as unemployed at January 01, 2004 or (ii) they have already signed a contract for old-age part-time employment before that date (Sozialgesetzbuch (SGB) VI §237)³¹.

³¹Mika and Krickl (2020) show that the majority of workers born in 1946 to 1951 used the “*Vertrauensschutz*” when retiring via the unemployment pathway, allowing early retirement at age 60.

1.A.3 Unemployment Insurance

In Germany, the UI system provides income replacement to eligible workers. Unemployed workers receive about 60% of their last net income as replacement payments and job search requirements are very low for older workers. Time spent on UI also increases future pension benefits. Older workers may therefore use UI as a stepping stone into retirement. Table 1.A.3 shows the maximum potential UI benefit duration at different ages. Over the observation period, the maximum potential benefit duration of older workers varies between 18 and 32 months, depending on workers' birth cohort.

Table 1.A.1: Requirements under Different Pathways

Requirements under Different Pathways	
Pathway	Requirements
Regular	At least 5 years of contributions No possibility of early retirement
Especially Long-term Insured	At least 45 years of contributions No possibility of early retirement
Long-term Insured	At least 35 years of contributions Possibility of early retirement (0.3% deductions for each month claiming before NRA)
Invalidity	At least 35 years of contributions <i>and</i> officially recognized disability of at least 50 % Possibility of early retirement (0.3% deduction for each month claiming before NRA)
Unemployment	At least 15 years of contributions <i>and</i> at least contributions in 8 out of 10 years before claiming <i>and</i> born before 1952 <i>and</i> (i) unemployed for at least 1 year after age 58 and 6 months <i>or</i> (ii) worked in old-age part-time employment for at least 24 months after age 55 Possibility of early retirement (0.3% deduction for each month claiming before NRA)

Notes: This table shows the NRA for the different pathways along with the ERA in brackets. A hyphen (-) indicates that a specific pathway is not applicable anymore.

1.B Data

1.B.1 Data Preparation

I start with a 2% random sample of all workers covered by social security who are identified in the IEB from 1975 to 2021. The data is provided at the spell level and each record includes a person ID, plant ID, start and end date of the respective spell, daily wage or benefit level as well as various personal characteristics and information on plant characteristics.

Since the wage information is generated from employer submitted employment records, the wage information in the data is highly reliable. However,

Table 1.A.2: Statutory Retirement Ages under Different Pathways

Statutory Retirement Ages under Different Pathways					
Cohort	Regular Pathway	Especially Long-term Insured Pathway	Long-term Insured Pathway	Invalidity Pathway	Unemployed Pathway
1943	65	-	65 (63)	62+1/12 (60)	65 (60)
1944	65	-	65 (63)	63 (60)	65 (60)
1945	65	-	65 (63)	63 (60)	65 (60)
1946	65	-	65 (63)	63 (60)	65 (60+1/12)
1947	65+1	-	65 (63)	63 (60)	65 (61+1/12)
1948	65+2	-	65 (63)	63 (60)	65 (62+1/12)
1949	65+3	-	65+1/3 (63)	63 (60)	65 (63)
1950	65+4	64	65+4 (63)	63 (60)	65 (63)
1951	65+5	63	65+5 (63)	63 (60)	65 (63)
1952	65+6	63	65+6 (63)	63+1/6 (60+1/6)	-
1953	65+7	63	65+7 (63)	63+7 (60+7)	-
1954	65+8	63+2	65+8 (63)	63+8 (60+8)	-

Notes: This table shows the NRA for the different pathways along with the ERA in brackets. A hyphen (-) indicates that the respective pathway has been abolished.

wages could be observed only up to the social security contribution ceiling, which implies that wages are right-censored. Right-censored wages are therefore imputed using a two-step procedure following Dustmann et al. (2009) and Card et al. (2013). In a procedure using “leave-one-out-means”, 270 tobit models are fitted separately by year, education group³² and region (east/west), including the following covariates: age, age², tenure, tenure², dummy for 20 or more employees, dummy for age above 40, interaction between dummy for age above 40 and age (age²) and dummy for women, respectively. Prices are deflated to 2015 prices using the consumer price index (CPI).

I follow Dauth and Eppelsheimer (2020) and create a yearly panel by taking one record for each individual using the spell that includes June 30th of the respective year.³³ If an individual has multiple employment spells, I keep the spell with the highest wage at the current employer. There are individuals who drop out of the data because they are not covered in the IEB anymore. There are several reasons for this: they could have started a civil service job, became self-employed, dropped out of the labor force, became unemployed (not registered), emigrated or died. In the analysis, I

³²I distinguish between three groups: (i) no qualification/qualification not recognized or missing, (ii) vocational training/high school diploma and (iii) university degree.

³³The sorting order of spells: 1) employment spells, 2) unemployment spells, 3) job seeker spells and 4) spells in subsidized employment and training measures.

Table 1.A.3: Maximum UI Benefit Duration at Different Ages (in Months)

Cohort	Age				
	55	57	59	61	63
1943	32	32	32	32	18
1944	26	32	32	32	18
1945	26	32	32	18	24
1946	26	32	32	18	24
1947	26	32	18	24	24
1948	26	32	18	24	24
1949	26	18	24	24	24
1950	26	18	24	24	24
1951	26	18	24	24	24
1952	26	18	24	24	24
1953	26	18	24	24	24
1954	26	18	24	24	24

Notes: This table shows the maximum UI benefit duration at different ages (in months). The minimum employment duration for new UI eligibility is 12. To receive 32 months of benefits, the minimum employment duration in the last 7 years is 64 months, for 26 months of benefits it is 54 months, for 18 months of benefits it is 36 months and for 24 months of benefits it is 48 months, respectively. *Source:* Gudgeon et al. (2019)

keep individuals in the sample until age 66 and assume zero earnings and employment for non-employed individuals.

1.B.2 Approximation of Pension Benefits

Unfortunately, the data does not include exact information on pension contributions and pension benefits. I therefore approximate those using information on workers' labor market biographies. To approximate pension benefits, I first calculate workers' pension contribution points and pension contribution years. In the data, I have information on all periods of employment and all periods receiving UI benefits, which are the main determinants of contribution periods³⁴. For each year, I identify the income from employment and unemployment that is relevant for the calculation of pension contributions. Note that labor earnings only count up to the social security limit and

³⁴In addition, I can identify periods of sick leave, as long as workers are in the labor market. I take into account pension contributions from sick leave, which count as pension contributions at 80% of last income.

Table 1.B.4: List of Occupations Excluded from the Analysis

Architects, garden architects
Brokers, landlords, agents, auctioneers
Entrepreneurs, managing directors, divisional managers
Management consultants, organizers
Chartered accountants, accountants, cost accountants, tax advisers
Members of parliament, ministers, elected officials, association leaders
Musicians
Artists' agents, visual or commercial artists, scenery or sign painters, artistic occupations
Interior, exhibition designers, window dressers
Photographers
Performers, professional sportsmen, auxiliary artistic occupations
Physicians, dentists, veterinary surgeons
Pharmacists
Non-medical practitioners, masseurs, physiotherapists and related occupations
Arbitrators, legal representatives
Commercial agents, travellers
University teachers, teachers

marginal employment was not subject to pension contributions for most of my observation period. Income from UI II benefits counted for 400 € (per month) from 2005 to 2006 and for 205 € from 2007 to 2010, while they do not count for pension contributions thereafter. Having identified this income, I divide this number by the average income of all insured to get the number of *contribution points* (for each worker in each year). Note that there is a maximum number of contribution points, as income only counts up to the social security ceiling. For example, a worker could have a maximum of 2.106 contribution points in 2021.

In addition to pension contribution points, I have to identify the *adjustment factor*, which penalizes early retirement and rewards retirement after the normal retirement age. For each worker, I identify the early retirement age at which he is eligible to claim a pension and the normal retirement age at which he can claim a pension without any deductions. If workers leave the data before they are first eligible to claim pension, I assume that they will claim pension at the earliest possible age. Note that I do not have information on disabilities and assume that no worker will enter retirement through the disability pathway.

Finally, I need information on the *pension value* to compute workers' pension. I use a sample average value of 27.1299 for the pension value.

Approximation of Left-censored Pension Contributions

In the data, observations are left-censored in 1975 and it is not possible to follow every worker since entering the labor market. As the main analysis focuses on workers born between 1943 and 1954, workers can be observed from the age of 21 or 32, depending on their birth cohort. Hence, I have to approximate pension contribution points and contribution years for earlier ages.

For the approximation, I use workers born between 1965 and 1975 who fulfill the same baseline restriction as workers in the main analysis sample and who I can observe since they entered the labor market. I then estimate the following regression to approximate the missing pension contribution points using the uncensored cohorts:

$$Y_{ia} = \sum_{k=17}^{35} \alpha_k \times 1[k = \text{Age}_{ia}] + \sum_{l=1}^9 \beta_l \times 1[l = \text{Occupation}_{ia}] + \sum_{m=1}^3 \gamma_m \times 1[m = \text{Education}_{ia}] \\ + \sum_{m=1}^3 \sum_{k=17}^{35} \delta_{m,l} \times 1[m = \text{Education}_{ia}] \times 1[k = \text{Age}_{ia}] + \epsilon_{ia},$$

where Y_{ia} represent the pension contribution points for worker i at age a . Occupation and education groups are defined as in the matching procedure (see section (1.4) and standard errors are clustered at the individual level³⁵.

I measure occupation and education in the pre-displacement year ($t - 1$) and approximate pension contribution points of censored cohorts using the following assumptions on the career-starting age: workers without education and those with vocational training enter the labor market at the age of 17 and those with a university degree at the age of 25. Approximated pension contribution points and contribution years are the same for each matched worker pair. I then use the approximated values until the first uncensored observation.

³⁵I do not control for workers' birth cohort, as contribution point are defined by a workers' income relative to the average income of all workers. This implicitly accounts for increasing wages over time.

1.C Appendix Figures and Tables

Table 1.C.5: Occupations of Displaced and Non-Displaced Workers

	Before Matching		After Matching	
	Displaced (1)	Controls (2)	Displaced (3)	Controls (4)
Resource extraction and production, agriculture	0.28 (0.45)	0.23 (0.42)	0.29 (0.45)	0.29 (0.45)
Construction	0.18 (0.39)	0.10 (0.30)	0.16 (0.37)	0.16 (0.37)
Mechanical, engineering and technical	0.26 (0.44)	0.30 (0.46)	0.28 (0.45)	0.28 (0.45)
Transportation, logistics, sales and services	0.19 (0.39)	0.23 (0.42)	0.19 (0.39)	0.19 (0.39)
Accounting, management and law	0.08 (0.28)	0.11 (0.31)	0.08 (0.27)	0.08 (0.27)
Arts	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Health, education and social affairs	0.00 (0.00)	0.03 (0.17)	0.00 (0.00)	0.00 (0.00)
Other	0.00 (0.00)	0.00 (0.02)	0.00 (0.00)	0.00 (0.00)
Observations	1467	11004	1220	1220

Notes: Statistics shown are means with standard deviations in parentheses. Occupations are measured in pre-displacement year $t - 1$. All workers fulfill the same baseline restrictions (see section 1.3.4).

Table 1.C.6: Lifetime Costs of Job Loss - Further Subgroups

	(1) PDV Income (incl. Estimated Pension Benefits)	(2) PDV Labor Earnings	(3) PDV Estimated Pension Benefits
<i>Panel A. Low wage</i>			
	-58046.8*** (19078.4)	-59418.8*** (18439.3)	-5693.3*** (1711.9)
Control Mean	542151.7	455013.7	81154.7
% of Mean	-10.71	-13.06	-7.02
Observations	620	620	620
<i>Panel B. Medium wage</i>			
	-84820.7*** (17404.6)	-83139.8*** (16940.9)	-7067.4*** (1548.4)
Control Mean	676134.2	566458.3	103868.0
% of Mean	-12.54	-14.68	-6.80
Observations	894	894	894
<i>Panel C. High wage</i>			
	-218861.3*** (30666.9)	-213873.4*** (29914.0)	-12369.2*** (2093.8)
Control Mean	1044965.7	903688.2	136493.7
% of Mean	-20.94	-23.67	-9.06
Observations	924	924	924
<i>Panel D. No/unrecognized Education</i>			
	-67680.6 (42257.7)	-72643.5* (41078.8)	-5941.0 (4435.9)
Control Mean	565263.9	469635.5	89936.1
% of Mean	-11.97	-15.47	-6.61
Observations	136	136	136
<i>Panel E. Vocational training/ high school diploma</i>			
	-113039.3*** (14109.6)	-110794.0*** (13371.0)	-8342.7*** (1421.8)
Control Mean	745340.5	629800.2	109985.4
% of Mean	-15.17	-17.59	-7.59
Observations	2184	2184	2184
<i>Panel F. University degree</i>			
	-451601.1*** (126442.5)	-448621.7*** (124408.5)	-12937.3** (5748.4)
Control Mean	1668213.7	1525619.2	138993.2
% of Mean	-27.07	-29.41	-9.31
Observations	120	120	120
<i>Panel G. No recession</i>			
	-132574.8*** (23141.8)	-129119.6*** (22105.2)	-9037.5*** (1993.8)
Control Mean	799878.7	686327.7	108417.9
% of Mean	-16.57	-18.81	-8.34
Observations	1170	1170	1170

(continued)

Life-time Costs of Job Loss - Further Subgroups (continued)

	(1)	(2)	(3)
	PDV Income (incl. Estimated Pension Benefits)	PDV Labor Earnings	PDV Estimated Pension Benefits
<i>Panel H. Recession</i>			
	-122174.8*** (21797.1)	-121746.6*** (20863.2)	-7879.4*** (1851.9)
Control Mean	763013.4	645216.5	112023.3
% of Mean	-16.01	-18.87	-7.03
Observations	1270	1270	1270
<i>Panel I. Resource extraction & production</i>			
	-80690.3*** (19617.4)	-79029.6*** (18691.8)	-6233.9*** (2170.5)
Control Mean	660486.2	551872.8	103611.7
% of Mean	-12.22	-14.32	-6.02
Observations	711	711	711
<i>Panel J. construction</i>			
	-33455.5 (25117.5)	-36481.9 (24188.6)	-2855.7 (2871.1)
Control Mean	596910.5	493017.2	96332.4
% of Mean	-5.60	-7.40	-2.96
Observations	390	390	390
<i>Panel K. Mechanics, engineering & technicians</i>			
	-223097.8*** (37559.1)	-217525.3*** (36366.0)	-13788.6*** (2550.8)
Control Mean	997026.6	867364.9	124992.0
% of Mean	-22.38	-25.08	-11.03
Observations N	686	686	686
<i>Panel L. Transportation, sales & services</i>			
	-99046.6*** (29946.6)	-99226.9*** (28026.9)	-7045.4** (3324.4)
Control Mean	693011.4	582455.5	105339.5
% of Mean	-14.29	-17.04	-6.69
Observations	464	464	464
<i>Panel M. Accounting, management & law</i>			
	-219180.2*** (68420.3)	-214667.9*** (65664.8)	-12032.2** (5650.5)
Control Mean	1044375.0	914935.5	122911.3
% of Mean	-20.99	-23.46	-9.79
Observations	188	188	188

Notes: This table shows the present discounted value (PDV) of income, labor earnings and estimated pension benefits for displaced workers relative to similar non-displaced workers by sub-groups. Sub-groups are defined in $t - 1$ (see section 1.5.2). Income includes labor earnings, UI benefits and estimated pension benefits. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, respectively.

Table 1.C.7: Effect of Job Loss on Retirement Probability - Late Career Job Loss

	(1) Pr(Retired)	(2) Estimated Monthly Pension Benefits
Age 55-59	0.0203** (0.00917)	
Age 60-64	0.0418** (0.0195)	-58.85*** (10.07)
Age 65-66	-0.00769 (0.0104)	-74.11*** (10.27)
Control Mean 55-59	0.0406	
Control Mean 60-64	0.428	1143.3
Control Mean 65-66	0.956	1201.8
# Cluster	650	650
Observations	15600	3908


Notes: Observations are at the worker \times age level. Standard errors (in parentheses) are clustered at the level of matched worker pairs. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, respectively.

Table 1.C.8: Effect of Job Loss on Retirement - Women

	(1) Pr(Retired)
Age 55-59	0.0166 (0.0186)
Age 60-64	-0.0331 (0.0299)
Age 65-66	0 (0.0111)
Control Mean 55-59	0.0939
Control Mean 60-64	0.640
Control Mean 65-66	0.974
# Cluster	326
Observations	7824

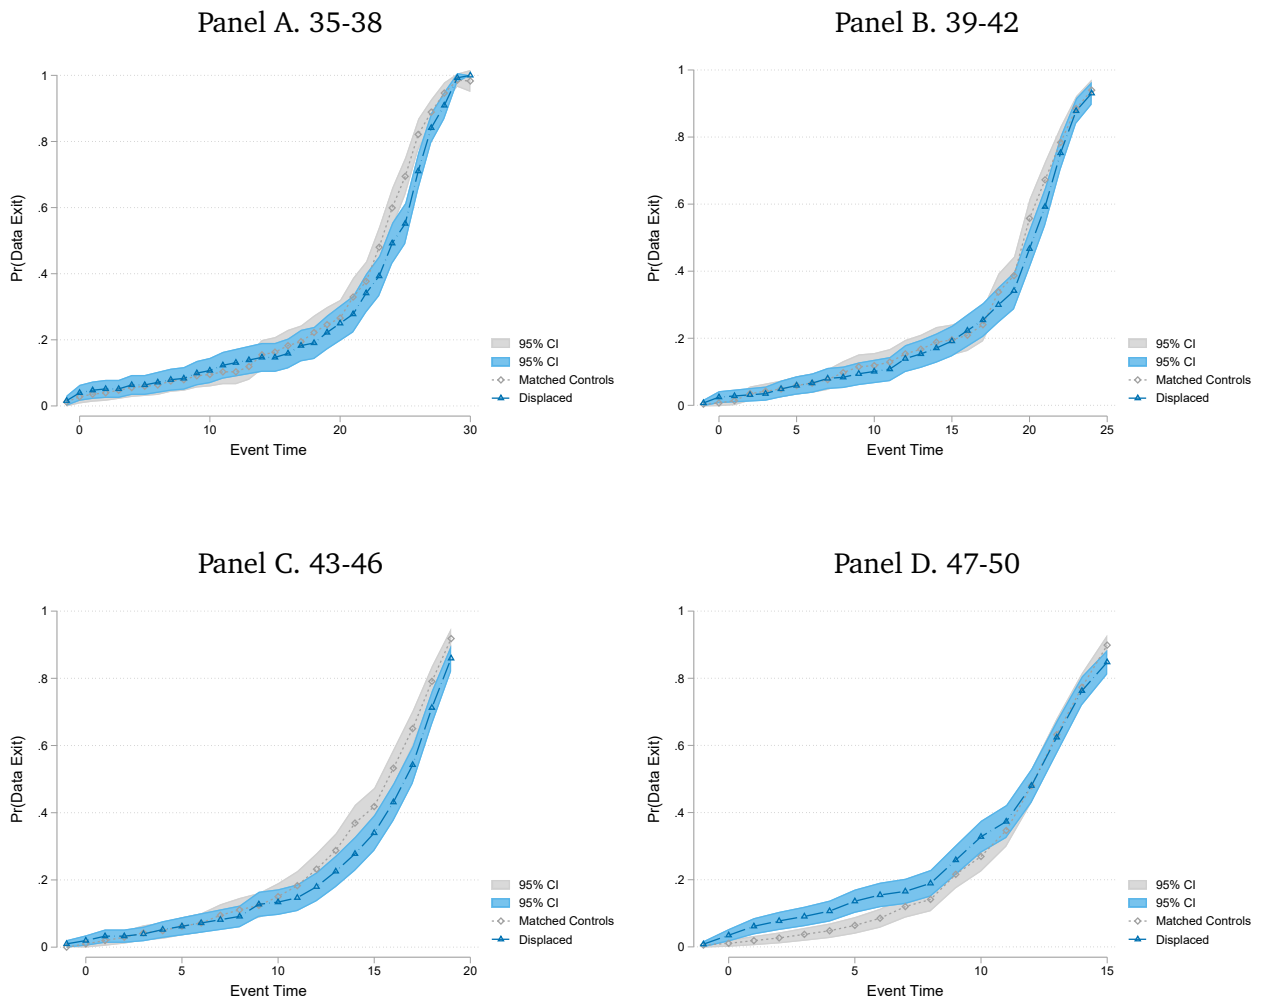
Notes: Observations are at the worker \times age level (from age 55-66). Standard errors (in parentheses) are clustered at the worker level (constant within matched worker pairs). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, respectively.

Figure 1.C.1: Descriptives Data Exit (Raw Means) - Displaced and Matched Control Workers



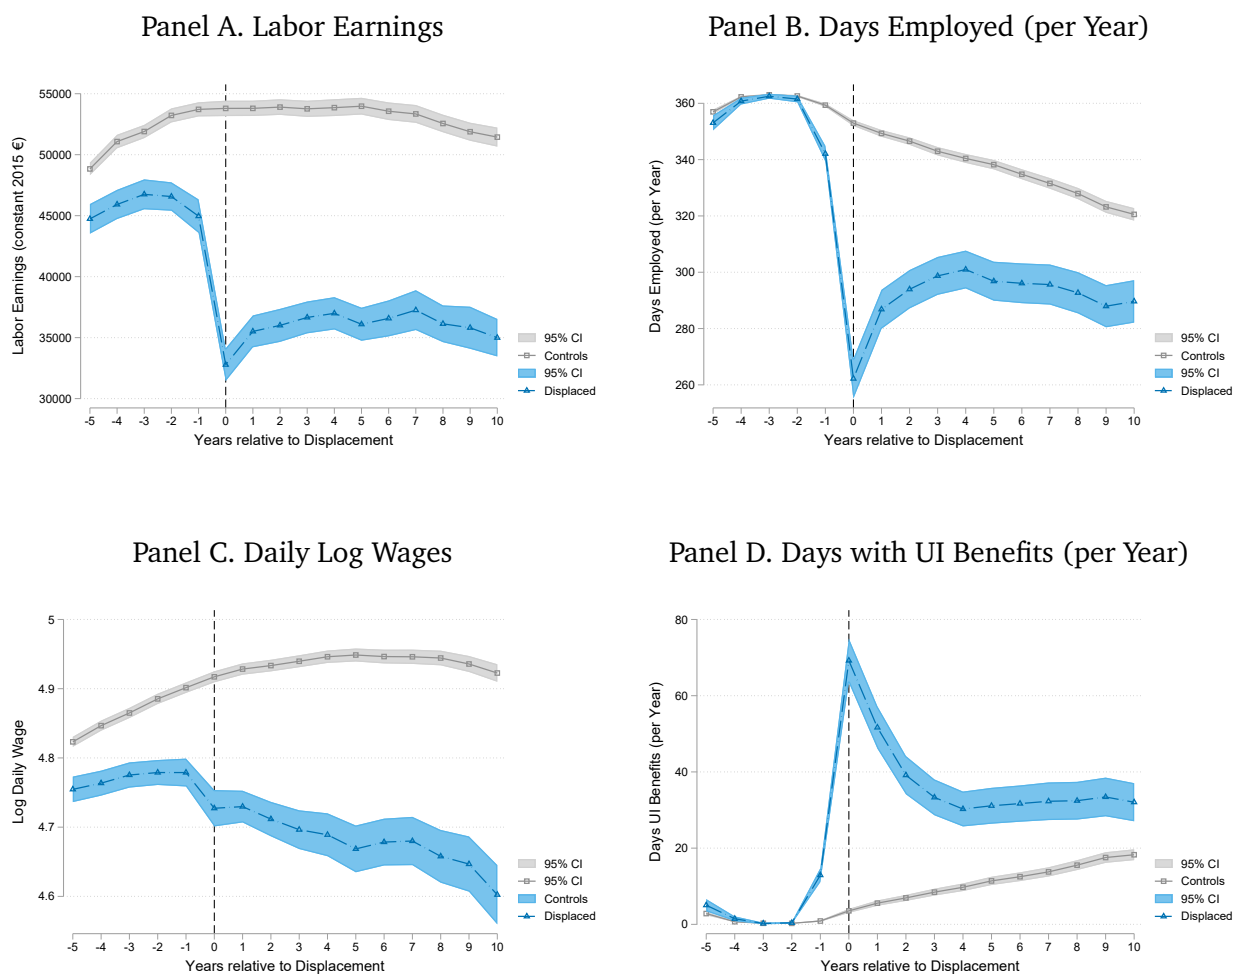
Notes: This figure shows raw means of the probability of data exit for displaced and matched control workers by event time.

Figure 1.C.2: Descriptives Data Exit (Raw Means) - Displaced and Matched Control Workers (by Age at Job Loss)



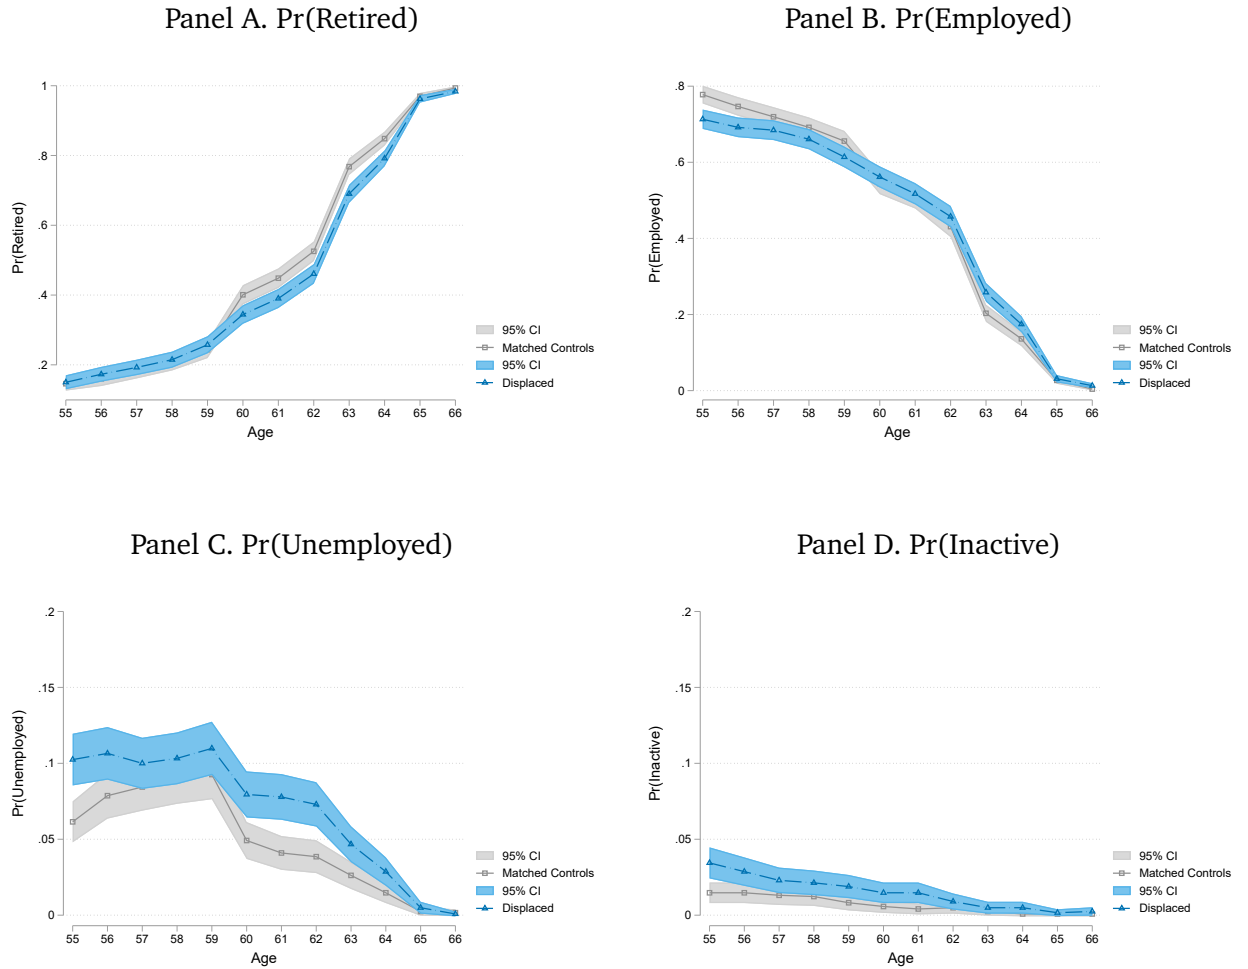
Notes: This figure shows raw means of the retirement probability for displaced and matched control workers by event time.

Figure 1.C.3: Displaced and Control Workers before and after Displacement - Raw Means before Matching




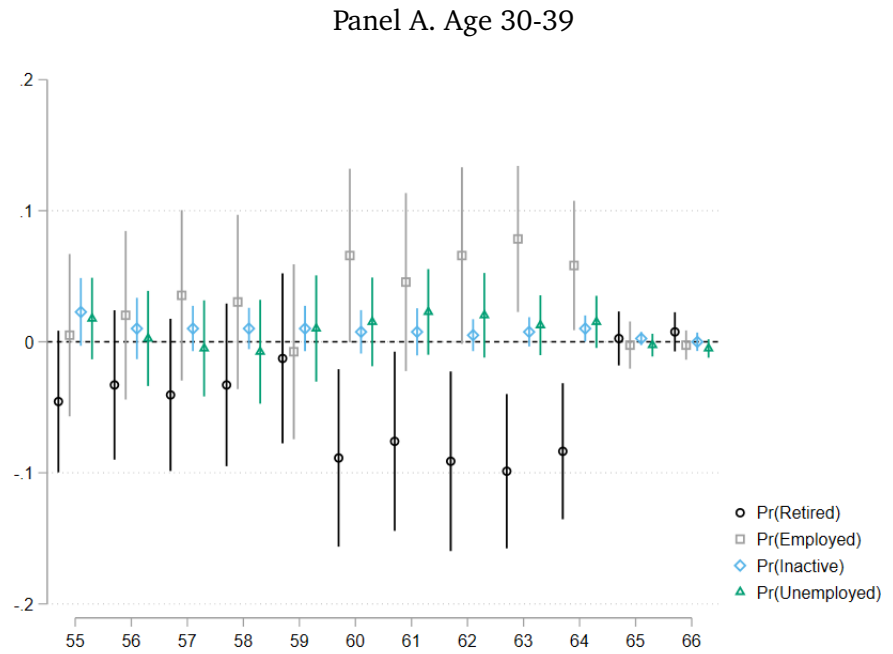
Notes: This figure shows raw means for displaced workers and unmatched control workers. Plant closures take place between June 30th in $t - 1$ and June 30th in t . The outcome data are measured in calendar years and observed from 1975-2021. Workers are displaced in any year from 1978-2004.

Figure 1.C.4: Pathways into Retirement (Raw Means) - Displaced and Matched Control Workers



Notes: This figure shows raw means for displaced workers and matched control workers.


Figure 1.C.5: Pathways into Retirement - By Age at Job Loss




(continued)

Pathways into Retirement - By Age at Job Loss (continued)

Panel C. Age 46-50



Notes: This figure shows the effect of job loss on the probability to be retired, employed, unemployed or inactive of displaced workers relative to matched control workers. Observations are at the worker \times age level. Coefficients are estimated using specification (1), separately by age group. 95% confidence intervals are derived from standard errors clustered at the level of matched worker pairs.

Figure 1.C.6: Effect of Job Loss on Retirement - Long-run

Notes: This figure shows the effect of job loss on the retirement probability of displaced workers relative to matched control workers. Observations are at the worker \times age level (from age 35-66). Coefficients are estimated using specification (1.1) (adding dummies for ages 35-54). 95% confidence intervals are derived from standard errors clustered at the level of matched worker pairs.

Moving to Opportunity - Together

joint work with Seema Jayachandran, Matthew Notowidigdo, Marie Paul,
Heather Sarsons and Elin Sundberg

Chapter Abstract

Many couples face a trade-off between advancing one spouse's career or the other's. We study this trade-off by analyzing the earnings effects of relocation and the effects of a job layoff on the probability of relocating using detailed administrative data from Germany and Sweden. Using an event-study analysis of couples moving across commuting zones, we find that relocation increases men's earnings more than women's, with strikingly similar patterns in Germany and Sweden. Using a sample of mass layoff events, we find that couples in both countries are more likely to relocate in response to the man being laid off compared to the woman. We then investigate whether these gendered patterns reflect men's higher earnings or a gender norm that prioritizes men's career advancement. To do this, we develop a model of household decision-making where households place more weight on the income earned by the man compared to the woman, and we test the model using the subset of couples where the man and woman have similar potential earnings. In both countries, we show that the estimated model can accurately reproduce the reduced-form results and can also quantitatively reproduce most of the observed female "child penalty."

This paper incorporates results from a previously-circulated working paper titled "Couples, Careers, and Spatial Mobility." We thank Eva Forslund, Daisy Lu, Angelo Marino, Isaac Norwich, and Nettie Silvernale for outstanding research assistance and Alessandra Voena for helpful comments. We also thank the German Institute for Employment Research (RDC-IAB) for generously providing the data and support with running programs remotely and the Economics and Business and Public Policy Research Fund at the University of Chicago Booth School of Business for financial support.

2.1 Introduction

Over the past half a century, women’s participation in the labor market has risen sharply in most OECD countries, and dual-earner couples have become the norm.¹ When each spouse contributes to household income, couples will have to make location decisions based on the potential job opportunities for each spouse. As a result, couples may face a trade-off: since job opportunities vary across regions, advancing one spouse’s career may come at the expense of the other’s, leading to the so-called “co-location problem” (Costa and Kahn, 2000).

Early models of the household predict that couples will make location decisions to maximize joint income (Mincer, 1978; Frank, 1978). Joint location decisions may therefore result in a gender earnings gap if men have higher earnings or higher earnings potential than women. Couples may choose to locate in areas that benefit the man’s career while the woman becomes the “trailing spouse”, working in a job that does not suit her skills or has lower earnings potential than if she was maximizing only her own earnings. However, numerous studies have shown that gender norms also influence household and individual decision-making (see, for example, Bertrand et al., 2015; Bursztyn et al., 2017).

In this paper, we use administrative data from Germany and Sweden to study the impact of moving on men’s and women’s earnings, and to test how much of the gender earnings gap from moving is due to differences in earnings potential versus gender norms. Using an event study design, we trace the earnings trajectories of heterosexual couples who move and find that moves disproportionately benefit men. While men’s earnings increase by about 11% and 5% in Germany and Sweden over the first five years following the move, women experience small changes in their earnings of 3% and -1%, respectively. These differences persist over the first 10 years following the move (men’s earnings increasing by about 17% and 11% in Germany and

¹For example, in 1970, 97% of German men and 47% of women aged 25-54 were in the labor force. By 2010, men’s labor force participation rate fell to 93%, while that of women increased to 81%, according to OECD statistics (<https://stats.oecd.org>). Also, in 2018, 65% of children aged 0-14 living in one-couple households had both parents working full-time and/or part-time in Germany, and this percentage goes up to 80% in Sweden (https://www.oecd.org/els/family/LMF_1_1_Children_in_households_employment_status.pdf).

Sweden, while women see a more modest increase of 10% and 7%).

We find that the earnings gap that emerges can be attributed to men experiencing an increase in wages, while women spend less time in the labor market, particularly in the first year following a move. We also find that the earnings gender gap following a move appears across all age groups and is most pronounced for couples in which both the man and the woman are between the ages of 20 and 29 at the time of the move.

We also study whether couples are more likely to move when the man is laid off as opposed to the woman. We use mass layoff events to generate plausibly exogenous job separations for both men and women in our sample of couples. In Germany, we find that the likelihood of moving increases following the layoff of either a man or a woman, but couples are nearly twice as likely to move when a man is laid off compared to when a woman is laid off. In Sweden, the likelihood that a couple moves doubles when the man is laid off, but does not change significantly when the woman is laid off. These results may help explain why women suffer larger earnings losses following a layoff relative to men: they are less able to take advantage of job opportunities in other localities (Illing et al., 2021).

To distinguish between different potential explanations for these reduced-form results, we consider a model of household decision-making in which households potentially place more weight on income earned by the man relative to the woman, as in Foged (2016). An intuitive prediction of the model is that in a standard unitary model of joint income maximization (net of migration costs), moves should not systematically benefit men in couples where the man and the woman have identical pre-move earnings and earnings potential. More generally, the gender gap in the effect of moves should be decreasing in the woman's share of household income and be reversed when the woman is the primary breadwinner. We find in both countries that the earnings gap that emerges following a move is indeed smaller among couples in which the woman has a higher predicted share of household income, consistent with potential earnings differences explaining some of the overall gender gap in the earnings effects of relocation. But even among couples where women have higher potential earnings, we find in both countries that men benefit more than women following a move.

With these empirical results as motivation, we then structurally estimate

the model parameters separately for each country. We test (and reject) the unitary model in both countries, with larger deviations from the unitary household benchmark in Germany than in Sweden. We also show that the model can reproduce the gender differences in the effects of a job layoff on the probability of moving, even though these results were not directly targeted in the model-based estimation. Lastly, we show that the model-based estimates can quantitatively account for most of the female “child penalty” in both countries by extending an existing model of the “child penalty” (Andresen and Nix, 2022) to allow households to place less weight on income earned by the woman compared to man’s (and calibrating this extended model using our country-specific estimates of the “discount” parameter β).

Our reduced-form results use a relatively standard event-study framework and mass layoff events to generate plausibly exogenous job separations. For both research designs, we present visual evidence that the identifying assumptions are plausible in both countries. Our model-based estimates require stronger assumptions, however. In particular, we assume that men and women have the same job opportunities and expected returns to migration conditional on predicted income. One way this assumption could be violated is if employers discriminate against women in making job offers to candidates in different commuting zones, perhaps in anticipation of women being less likely to be able to accept offers to relocate. To address this concern, we have replicated our heterogeneity analysis by female share of predicted income using different prediction models that allow for gender discrimination, and we find broadly similar results.

Overall, we conclude that our empirical results and model-based estimates suggest that a gender norm that prioritizes men’s career advancement can simultaneously (and parsimoniously) account for three distinct gender differences in labor market outcomes: the earnings effects of relocation, the probability of moving following a job layoff, and the earnings effects of the birth of a child.

Our paper relates to a large literature on the source of gender gaps in labor market outcomes. A number of papers have found that child penalties play an important role in the remaining gender gap (Angelov et al., 2016; Cortes and Pan, 2022; Kleven et al., 2019b,a). Women, who typically take over more care responsibilities than men, have disadvantages when long working hours or

working particular hours is rewarded (Bolotnyy and Emanuel, 2022; Goldin, 2014). Women also show a lower willingness to commute (Le Barbanchon et al., 2020). In addition, social norms or psychological attributes such as being willing to compete, risk preferences, and self-confidence may directly affect job search and wages (e.g. Bertrand et al., 2015; Buser et al., 2014; Cortes et al., 2021; Wiswall and Zafar, 2017). A further potential explanation, which is the focus of this paper, is that married women may take less advantage of career enhancing long-distance moves or may even experience earnings losses as a tied mover.

In this space, a number of papers have examined joint location decisions and the rise of female labor force participation. Early papers, such as Mincer (1978), model household decision-making under the constraint that, within a couple, one individual is typically “tied”. That is, the individual benefits less from migration made under household decisions than if they could move individually. These early papers document women’s increased labor force participation as a constraint on individual optimization, but do not directly test how migration decisions are made. A number of papers have since empirically documented couples’ location decisions, noting that married couples are less likely to move than single individuals, and also move to different areas (Costa and Kahn, 2000; Compton and Pollak, 2007; Rabe, 2009; Blackburn, 2010a). Studies that attempt to directly study the impact of moving on gender inequality have typically had to use a selected sample or are unable to establish causality. For example, Burke and Miller (2018) use military spouses to estimate the impact of an exogenous move on the spouse’s labor market outcomes, and Nivalainen (2004) looks at families in Finland and shows that most moves occur to help the man’s career. By using administrative data from Sweden and Germany and an event study design, we contribute to this literature by estimating the causal impact of couples moving on men’s and women’s earnings covering a large and fairly representative sample of heterosexual couples in the entire working-age population.

Our paper also relates to more recent research examining the implications of location decisions on gender inequality. Fadlon et al. (2022) examine how early labor market choices impact career and family outcomes for male and female physicians in Denmark. Exploiting the lottery system that allocates physicians to initial internships, the authors find that the geographic location

of the internship explains a large fraction of gender inequality in human capital accumulation and wages, suggesting that women may be more tied to location. Venator (2020) uses the NLSY97 to test how unemployment insurance generosity affects couples' migration decisions, finding that access to UI increases migration rates as well as women's post-move earnings. Relative to this work, we develop and test an alternative model-based explanation that allows for a gender norm that prioritizes the man's career within the couple.

The remainder of the paper proceeds as follows. We describe the two administrative datasets as well as our sample and variable construction in section 2.2. Section 2.3 describes our empirical strategy and results are presented in section 2.4. Section 2.5 develops a model of household decision-making and presents additional empirical results motivated by the model. Section 2.6 provides additional evidence for gender norms. Finally, section 2.7 concludes.

2.2 Data

We use administrative data from Sweden and Germany to test whether moves disproportionately benefit men in heterosexual couples. These datasets are ideal for three reasons. First, in each dataset, we have geographic information on the place of residence for each spouse that is necessary to investigate the effects of joint moves. Second, the data include detailed labor market histories of both spouses, allowing us to precisely account for spouses' pre-move employment outcomes and study the post-move dynamics. Third, we can identify mass layoff events at the establishment level, using them as an exogenous negative labor market shock that could lead to a move.

2.2.1 German Data

For Germany, we use a 25% random sample of married couples that can be identified in the administrative data base Integrated Employment Biographies (IEB).² The IEB includes all employees subject to social security (this excludes

²The data product we use is produced by the Institute of Employment Research (IAB). The data are processed and kept by the IAB according to Social Code III. The data contain sensitive information and are therefore subject to the confidentiality regulations of the German Social Code (Book I, Section 35, Paragraph 1). The data are held by the IAB, Regensburger St 104,

civil servants and self-employed), all people who receive unemployment benefits, and those who have been registered as searching for a job. Married couples are identified according to the method of Goldschmidt et al. (2017): for two people to be matched as a couple, the spouses have to live at the same location, have a matching last name, be of different sexes, and have an age difference of less than 15 years.³ The identification of couples is done every year on June 30 which implies that our data only includes couples of whom both spouses have a record in the IEB for that particular date. It also means that we cannot be sure whether two individuals remain a couple in cases in which at least one of the two individuals does not have a data record on June 30th in the following years.

The dataset consists of day-to-day information on every period in employment covered by social security, every period of receiving unemployment insurance benefits, as well as information on periods of job search and participation in subsidized employment and training measures. For each period, it contains information on the corresponding wages and benefit levels. The wage information is accurate, as the employer has to report wages for social security purposes. In addition, the data include a rich set of personal characteristics such as occupation, nationality, year of birth, education, and job requirement level. For each employee, we also observe information on the employers, such as firm size, average wage at the firm and industry, obtained from the Establishment History Panel (BHP)⁴. In our analysis, we use this link between employees and firms to identify mass layoffs.

2.2.2 Swedish Data

We use individual-level administrative data from Sweden from the GEO-Sweden database. The database covers the entire Swedish population of 10 million people, whom we can track over time starting in 1990. In addition, we can identify the building in which individuals reside, allowing us to identify couples. Specifically, we identify heterosexual couples as individuals of the

D-490478 Nuremberg, email: iabiab.de, phone: +49/911 1790. If you wish to access the data for replication purposes, please get in contact with the authors and the IAB.

³This identification method increases the likelihood of identifying certain types of couples: 1) older couples, 2) more conservative couples, and 3) couples living in smaller buildings (Goldschmidt et al., 2017).

⁴Throughout this paper, we use the term firm for simplicity. Note that we can only identify establishments and are unable to link them to firms.

opposite sex who move to and from the same building in the same year. We restrict the data in several ways to construct our final sample of couples, described in detail in sub-section 2.2.4.

2.2.3 Moving Across Commuting Zones

To focus on couples that change local labor markets when they relocate, we study moves across commuting zones using district-level information on each couple's place of residence. Kosfeld and Werner (2012) define commuting zones in Germany as districts connected through high commuter flows and identify 141 commuting zones in Germany. For Sweden, we use Statistics Sweden's concept of *FA-regioner* to identify 60 commuting zones⁵, see Figure 2.1.

In the German data, the information on the place of residence is only determined at the end of each year for most spells⁶. We therefore allow for the possibility that one spouse moves in year t while the other follows in year $t + 1$.

⁵More details here: https://www.scb.se/contentassets/1e02934987424259b730c5e9a82f7e74/fa_karta.pdf.

⁶For employment spells (BeH), which form the bulk of observations, the information on the place of residence is determined at the end of each year. For job seeker spells (ASU), unemployment benefit spells (LeH), and participant in training measures spells (MTH and XMTH), the information on the place of residence applies to the beginning of the original period. Only for unemployment benefit II recipient spells (LHG) and XASU spells (ASU spells reported by municipal institutions) the information applies to the entire period of original observation.

2.2.4 Sample Selection, Variable Definition, and Descriptive Statistics

Movers Sample

In our analysis, we consider all joint moves of couples occurring between 2002-2007 Sweden and 2001-2011 in Germany. During the observation period, a few couples experienced multiple long-distance moves. We consider only their first move, because future outcomes may be influenced by the first move. We therefore abstract from repeated migration.

We exclude couples where neither spouse is 25 to 45 years old at the time of the move, as well as couples with an age difference larger than 15 years.⁷ In the Swedish data, we use the receipt of student benefits to identify and exclude couples in which at least one person is a student in the five years preceding a move. In the German data, we are excluding couples in which at least one person is in education (e.g. apprentice, intern) in the five years before a move.⁸ Finally, couple-years in which one spouse is above 60 or below 16 years old are excluded.

We construct a panel that includes all couples that we observe at least 2 years before the move to 4 years thereafter (i.e., a partially balanced panel). Our final sample consists of 12, 747 moving couples in Germany and 44, 499 couples in Sweden.

Variable Definitions and Descriptive Statistics

The main outcome variable that we consider in our analysis is gross yearly wage income (in 2017 euros) of each spouse. For non-working spouses, the wage income is zero. Changes in wage income may therefore be either due to changes at the extensive or intensive margin.

Table 2.1 presents descriptive statistics for our two samples. The average age of couples is similar in each sample. Education levels are different, in large part due to differences in the education systems. Sweden has a lower part-time employment rate for women. For the age group from 25 to 54 years old, in 2010, the share of part-time workers for men and women were 5.6

⁷We do this to ensure that we do not accidentally pick up on child-parent pairs.

⁸We exclude students so that any income changes following a move are not due to initial entry into the labor market.

and 39.1% in Germany, and 5.0 and 13.4% in Sweden⁹. Table 2.2 presents descriptive statistics for the layoffs samples.

Table 2.1: Summary Statistics for Movers Sample

	Germany		Sweden	
	Men (1)	Women (2)	Men (3)	Women (4)
Age	36.16 (6.17)	33.87 (6.12)	35.00 (6.86)	32.71 (6.33)
Compulsory schooling	0.01 (0.11)	0.03 (0.17)	0.13 (0.33)	0.12 (0.33)
High school	0.05 (0.21)	0.06 (0.25)	0.48 (0.50)	0.44 (0.50)
Vocational training	0.60 (0.49)	0.68 (0.46)	0.07 (0.26)	0.04 (0.20)
Some college	0.34 (0.47)	0.22 (0.42)	0.57 (0.50)	0.56 (0.50)
Potential experience	17.17 (6.43)	15.17 (6.32)	15.32 (0.43)	13.03 (0.45)
Wage income (1000s EUR)	44.11 (39.95)	19.79 (22.04)	28.94 (19.48)	16.60 (14.05)
Employed	0.88 (0.33)	0.78 (0.41)	0.89 (0.31)	0.84 (0.36)
Unemp. benefits (1000s EUR)	0.61 (2.06)	0.39 (1.40)	0.90 (2.72)	0.99 (2.59)
Days receiving UI benefits (per year)	20.80 (66.30)	20.92 (70.20)	23.94 (64.71)	24.51 (62.95)
At least 1 child	0.62 (0.49)	0.62 (0.49)	0.66 (0.47)	0.66 (0.47)
Non-native	0.08 (0.27)	0.08 (0.28)	0.13 (0.34)	0.14 (0.35)
Observations	12747	12747	44499	44499

Notes: This table displays means and standard deviations (in parentheses) for different outcomes in the period before the move ($t - 1$) in Germany and Sweden for the movers sample.

⁹These statistics are from OECD's indicator of share of employed in part-time employment, by sex and age group (<https://stats.oecd.org/>). If we consider the Swedish definition of part-time employment –less than 35 hours a week, as opposed to OECD's definition of less than 30 hours–, we find a part-time employment rate of about 30% for Sweden in their own statistics (<https://pxweb.nordicstatistics.org>).

Table 2.2: Summary Statistics for Job Layoffs Sample

	Germany				Sweden			
	Layoff Men		Layoff Women		Layoff Men		Layoff Women	
	Men (1)	Women (2)	Men (3)	Women (4)	Men (5)	Women (6)	Men (7)	Women (8)
Age	38.26 (4.85)	36.51 (5.65)	40.77 (5.93)	38.24 (5.03)	36.49 (4.95)	34.89 (5.63)	39.02 (6.21)	36.41 (5.06)
Compulsory schooling	0.01 (0.09)	0.02 (0.15)	0.01 (0.08)	0.01 (0.10)	0.11 (0.31)	0.08 (0.28)	0.15 (0.36)	0.09 (0.29)
High school	0.09 (0.28)	0.09 (0.28)	0.06 (0.23)	0.09 (0.29)	0.55 (0.50)	0.53 (0.50)	0.54 (0.50)	0.55 (0.50)
Vocational training	0.76 (0.43)	0.79 (0.41)	0.78 (0.41)	0.79 (0.41)	0.12 (0.32)	0.03 (0.18)	0.06 (0.24)	0.04 (0.20)
Some college	0.14 (0.35)	0.10 (0.30)	0.15 (0.36)	0.11 (0.31)	0.06 (0.23)	0.15 (0.36)	0.06 (0.25)	0.11 (0.31)
Potential experience	19.76 (5.07)	18.07 (5.85)	22.03 (6.13)	19.69 (5.29)	17.08 (5.32)	15.49 (5.88)	19.86 (6.73)	17.13 (5.53)
Wage income (1000s EUR)	43.87 (27.14)	16.95 (17.11)	41.14 (30.97)	28.17 (15.63)	38.21 (16.08)	17.39 (13.47)	32.71 (19.12)	25.43 (12.16)
Employed	1.00 (0.00)	0.85 (0.36)	0.93 (0.25)	1.00 (0.00)	1.00 (0.00)	0.90 (0.30)	0.93 (0.25)	1.00 (0.00)
Unemp. benefits (1000s EUR)	0.60 (1.61)	0.30 (1.19)	0.43 (1.70)	0.45 (1.21)	0.44 (1.72)	0.73 (2.16)	0.54 (2.16)	0.47 (1.64)
Days receiving UI benefits (per year)	16.47 (41.26)	18.60 (70.34)	16.16 (59.97)	18.77 (46.28)	11.49 (39.25)	15.32 (48.20)	11.88 (46.87)	9.86 (35.75)
At least 1 child	0.57 (0.49)	0.57 (0.49)	0.47 (0.50)	0.47 (0.50)	0.91 (0.28)	0.92 (0.28)	0.89 (0.31)	0.90 (0.31)
Non-native	0.08 (0.28)	0.07 (0.26)	0.05 (0.23)	0.05 (0.22)	0.10 (0.30)	0.11 (0.31)	0.10 (0.30)	0.09 (0.29)
Observations	6828	6828	4458	4458	8052	8052	6768	6768

Notes: This table displays means and standard deviations (in parentheses) for different outcomes in the period before the layoff ($t - 1$) in Germany and Sweden for the job layoffs sample.

2.3 Empirical Strategy

We follow an event study approach to estimate the impact of a move on men's and women's labor market outcomes. The usual identification assumptions for event an event study design are no-anticipation and parallel trends. This involves that the event (moving) is not determined by the outcome (earnings or employment). In our setting, it is likely that individuals choose to move in response to income or employment shocks. However, we are particularly interested in whether couples are equally likely to move in response to a shock to a man's or a woman's career. The biggest threat to our strategy is that couples move when women (or men) are choosing to exit or enter the labor market, or to work less. For example, if couples choose to move when they are starting a family, the move will coincide with women temporarily leaving the labor market. We therefore control for an individual's potential experience and education level, as well as calendar year and child event-time indicators. These controls, and the fact that we exclude students, should account for potentially endogenous reasons why couples might move.

Our main estimation equation is

$$\begin{aligned}
 Y_{ist}^g = & \sum_{j \neq -1} \alpha_j^g \times 1[j = t] + \sum_k \beta_k^g \times 1[k = \hat{\text{exp}}] + \sum_p \gamma_p^g \times 1[p = \text{educ}_{is}] \\
 & + \sum_y \nu_y^g \times 1[y = s] + \sum_m \tau_m^g \times 1[m = t_{ch}] + \theta_n^p \times X + \epsilon_{ist}^g \quad (2.1)
 \end{aligned}$$

where the outcome of interest is individual i 's wage income in year s and event time t . The first term consists of event-time indicators, which we estimate for five years before and ten years after a move. We estimate equation 2.1 separately by gender g and include controls for potential experience ($\hat{\text{exp}}$), education level (educ), calendar year ($y = s$), and child event-time ($m = t_{ch}$).¹⁰ Standard errors are clustered at the individual level.

¹⁰There are five education levels: compulsory schooling, high school, vocational training, some college, and college.

2.4 Results


2.4.1 Descriptive Results

We begin by separately plotting men's and women's unconditional wage income and employment status following a move, shown in Figure 2.2. Panels (a) and (b) show the wage income for German and Swedish couples who move together for the first time. Both men's and women's incomes are relatively flat prior to the move in time 0, after which men's income steadily increases. For both countries, we see a slight dip in women's earnings around the time of a move followed by steady income growth.


These moves partly appear to occur following a period of unemployment. Panel (c) and (d) show that men and women receive fewer days of unemployment benefits following a move, although there is a spike in benefit collection for women in the year and or the year after a move. These results provide initial evidence that these moves may be for the benefit of men's careers.

Figure 2.1: Maps of Commuting Zones

(a) Germany

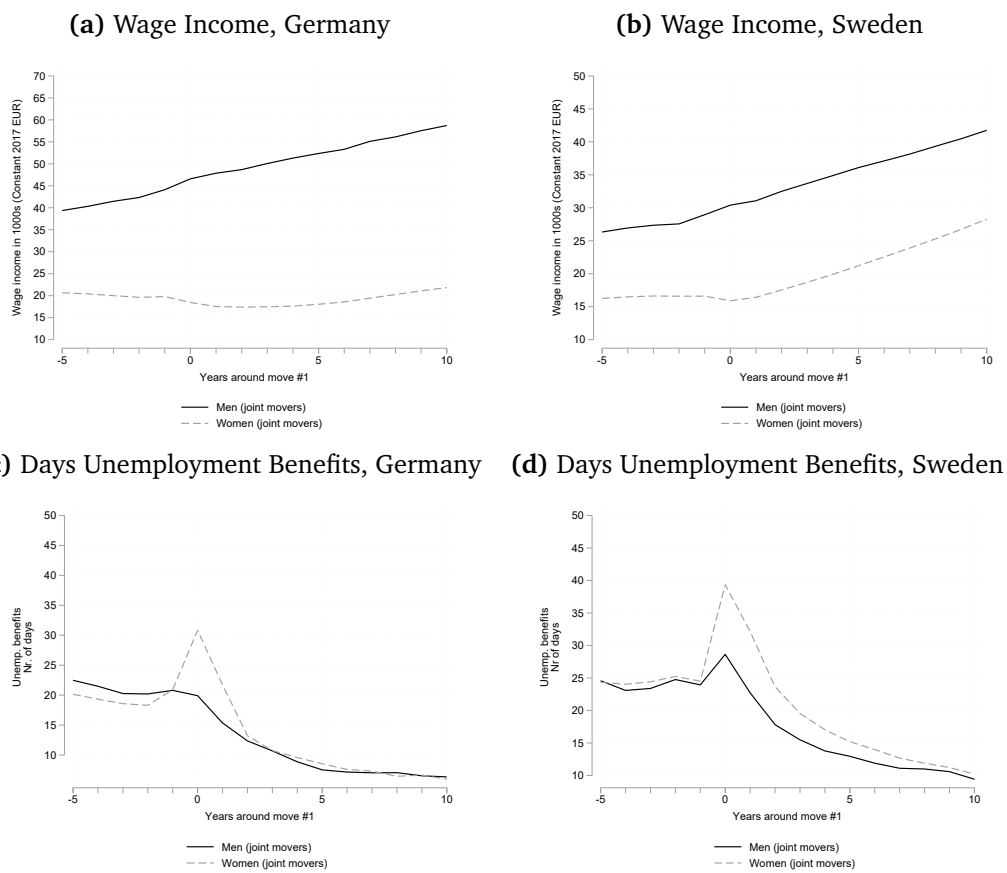


(b) Sweden



Notes: This figure displays the maps of the commuting zones in Germany and Sweden. Commuting zones in Germany follow Kosfeld and Werner (2012).

Figure 2.2: Relationship between Moving and Labor Earnings and Employment



Notes: This figure displays means for different variables in each country from $t - 5$ to $t + 10$ relative to the first move, per gender.

2.4.2 Main Results: Earnings Effects of Moving Across Commuting Zones

We now turn to our main estimation strategy, in which we compare the labor market outcomes for men and women who move while controlling for experience, education, calendar year, and child event-time indicators. We plot the coefficients from estimating equation 2.1 in Figure 2.3. The coefficients are plotted relative to the average of the outcome variable in the year before the move ($t - 1$).

In both Germany and Sweden, a gap between men's and women's earnings emerges the year of the move and steadily grows over time. Five years after a move, men are earning about €8,000 and €3,000 more than they were in the year prior to the move, while women are earning about €2,000 and €1,000 more in Germany and Sweden respectively.

To investigate whether spouses' earnings responses are driven by changes in employment or in wages, panels (c) and (d) of Figure 2.3 and (a)-(f) of Figure 2.B.2 show the effects of a move on various employment measures of men and women. In Germany, the number of days a person is employed increases by 20 days per year in the year immediately following a move for men and by less than 10 days per years for women. However, employed days continue to increase over time and eventually converge. We also see a spike in the number of days an individual collects unemployment benefits in year following a move that is much more pronounced for women than for men (17 days versus 7 days). These results suggest that at least part of the divergence in men's and women's earnings is due to women leaving employment for a period of time following a move.


The results in Figure 2.3 indicate that relocation increases wage earnings of men more than women in absolute terms, and Figure 2.4 indicates that this is true in proportional terms, as well. Figure 2.4 normalizes the event study estimates in Figure 2.3 (panels (a) and (b)) by the average income of men and women in each country in the year prior to the move.¹¹ These results show that moving increases the average earnings growth for men by a greater percentage than women; specifically, 10 years after the move,

¹¹This normalization follows the approach in the recent "child penalty" literature (see, e.g. Kleven et al. (2019b))


men experience a 9.6 percentage point higher earnings growth compared to women in Germany, and in Sweden the gender gap is 4.3 percentage points. Interestingly, in both countries men and women experience long-run increases in earnings, but men experience greater earnings growth in both absolute and percentage terms. The fact that average earnings increase significantly for both members of the household is consistent with non-negligible migration costs.

Figure 2.3: Impact of Move on Labor Earnings and Employment


(a) Wage Income, Germany




(b) Wage Income, Sweden



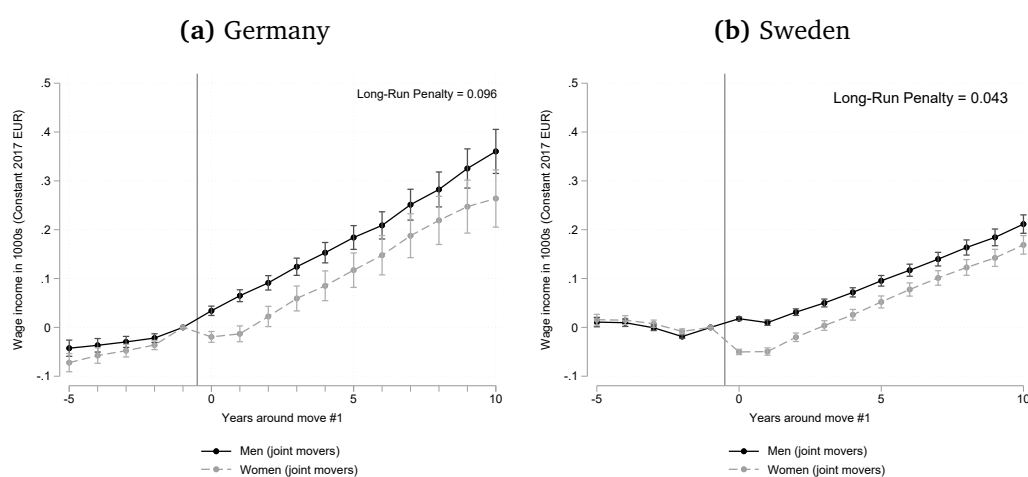
(c) Days Unemployment Benefits, Germany



(d) Days Unemployment Benefits, Sweden



Notes: This figure displays the event study results that estimate the effect of moving on different outcomes in each year relative to the year before the move ($t - 1$). Each point estimate has a corresponding 95% confidence interval calculated using standard errors clustered at the individual level. The regressions are run separately by gender. The coefficients and standard errors (in parentheses) in the upper right corner of each figure are 6 and 11-year averages of the post-move point estimates (from $t = 0$ to $t = 5$ and $t = 10$), in this order, for men (M) and women (W).

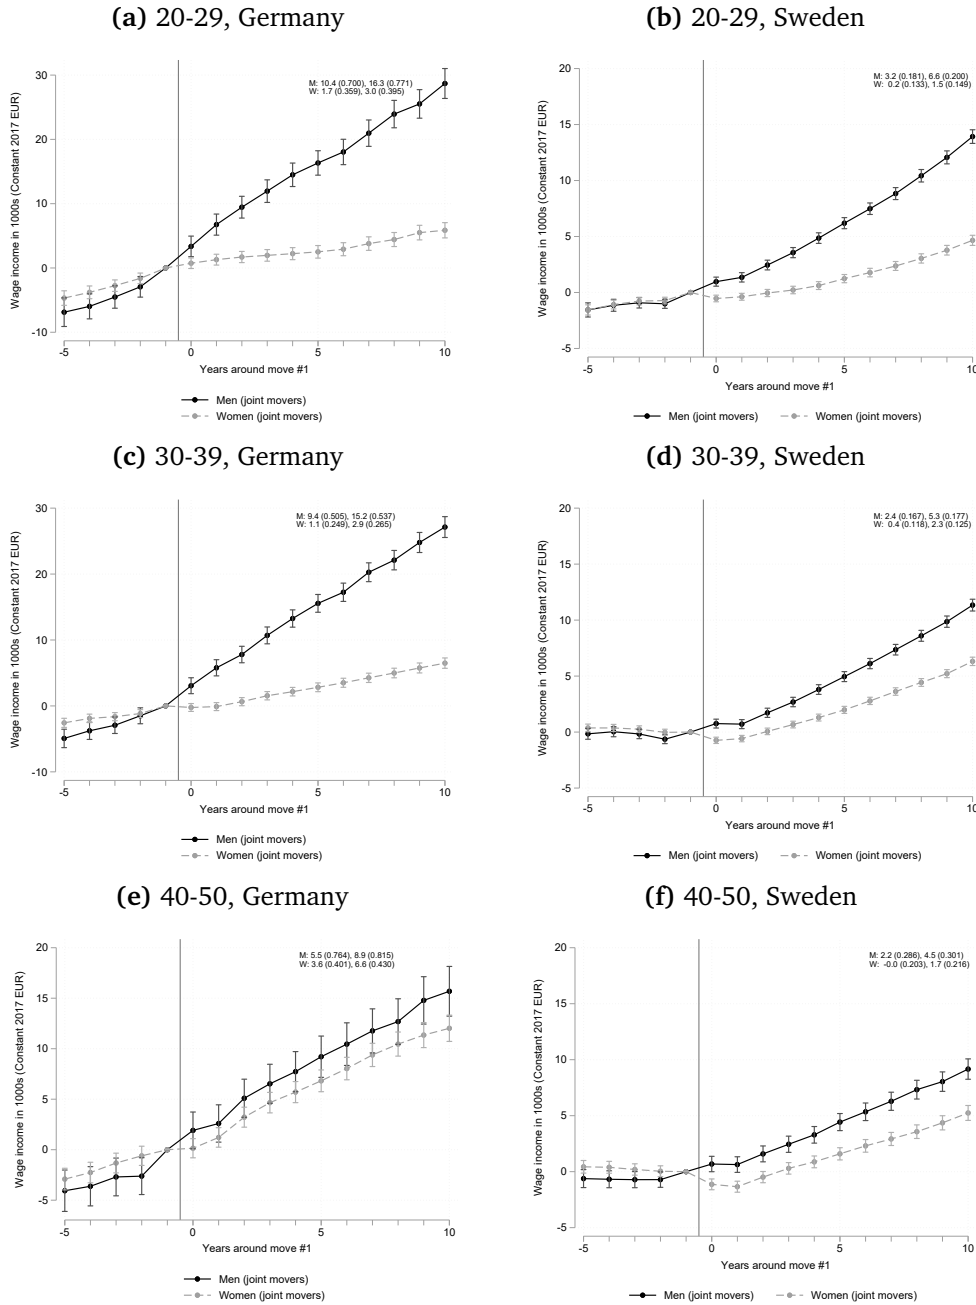
Figure 2.4: Proportional Impact of Move on Wage Income

Notes: This figure displays the event study results that estimate the proportional effect of moving on wage income in each year relative to the year before the move ($t - 1$). Each point estimate has a corresponding 95% confidence interval calculated using standard errors clustered at the individual level. The regressions are run separately by gender. The long-run penalty is calculated as in Kleven et al. (2019b) and it measures the percentage by which women are falling behind men due to move at event time $t = 10$.

2.4.3 Heterogeneity

Previous research showed that young individuals are more likely to move (Polachek and Horvath, 2012) and that the returns to moving are larger for younger individuals (Bartel, 1979). To test whether the treatment effects vary with respect to spouses' age, we define age groups based on the average of the spouses (in pre-move year $t - 1$). We define age groups for the following age intervals: 20 – 29, 30 – 39, and 40 – 50. The results, displayed in Figure 2.5, show that the returns to moving decline with increasing age. For both spouses, the average treatment effects on wage income are the largest for younger couples and the lowest for older couples. We see gender differences in the returns to moving for all age groups, but they are smallest among the oldest age group, where men's returns are relatively low.

Figure 2.5: Impact of Move on Wage Income – By Age Groups



Notes: This figure displays the event study results that estimate the effect of moving on wage income in each year relative to the year before the move ($t - 1$) for different age groups. Each point estimate has a corresponding 95% confidence interval calculated using standard errors clustered at the individual level. The regressions are run separately by gender. The coefficients and standard errors (in parentheses) in the upper right corner of each figure are 6 and 11-year averages of the post-move point estimates (from $t = 0$ to $t = 5$ and $t = 10$), in this order, for men (M) and women (W).

2.4.4 Mass Layoff Results

The previous results show the emergence of a significant earnings gap following a joint move, with men seeing more earnings growth following a move than women. In this section, we use mass layoff events to test whether couples are equally likely to move for men's and women's careers following a layoff.

We restrict our sample to the set of couples in which one person in the couple loses his or her job as part of a mass layoff. We define a mass layoff as a reduction in a firm's workforce by more than 30%. We exclude workplaces with fewer than 50 employees, as well as firms where 30% or more of employees jointly move to another workplace.¹² For the sample of mass layoff movers, the same age and student restrictions are imposed as described in section 2.2. In addition, we restrict the sample to individuals who have earnings of at least €8,000 in the year before the mass layoff occurs. We further focus on individuals who have worked at the firm at which they are laid off for at least one year, to minimize the possibility that we are picking up on temporary workers. We again consider an individual's first layoff.

We show descriptively how men's and women's earnings and employment change following a mass layoff in Figure 2.B.8. For both men and women, wage income drops sharply the period of the mass layoff ($t = 0$). Men's income appears to recover to its $t = -1$ level about five years after the layoff whereas for women the recovery is slower (panels a and b).

In Table 2.3 we examine how the likelihood of moving depends on whether a man or a woman within a couple is laid off. We regress an indicator that takes the value one if a couple moves in the year of a mass layoff (or the year after) on indicators for either the man or the woman being laid off. Column 1 shows that, for Germany, the likelihood of moving increases by 0.7 percentage points when a man is laid off (relative to a baseline moving rate of 0.7%) and by 0.1 percentage points when a woman is laid off. These estimates do not change when we include age and commuting zone fixed effects (columns 2 and 3) and the pattern is similar for Sweden.

¹²We assume that in this case, the firm has been acquired or has split of part of its operations.

Table 2.3: Impact of Layoffs on Moving Probability

	Germany			Sweden		
	(1)	(2)	(3)	(4)	(5)	(6)
Layoff Men	0.00680*** (0.00147)	0.00581*** (0.00149)	0.00594*** (0.00148)	0.0147*** (0.00190)	0.0148*** (0.00189)	0.0146*** (0.00189)
Layoff Women	0.00127 (0.00143)	0.00149 (0.00143)	0.00174 (0.00144)	-0.00294** (0.00133)	-0.000448 (0.00133)	-0.000577 (0.00134)
Age FE		✓	✓		✓	✓
CZ FE			✓			✓
# Layoff Men	6828	6828	6828	8052	8052	8052
# Layoff Women	4458	4458	4458	6731	6731	6731
Mean	0.00719	0.00719	0.00719	0.0153	0.0153	0.0153
P-Value	0.041	0.191	0.146	<0.001	<0.001	<0.001
Observations	165449	165449	165449	263680	263680	263680

Notes: This table displays point estimates and standard errors clustered at the individual level (in parentheses) for the impact of layoffs for men and women on the probability of moving in t or $t + 1$. The p-values refer to the test of whether the men and women layoff coefficients are equal. These regressions are run on the full sample of couples
* $p < .1$, ** $p < .05$, *** $p < .01$

2.5 Model-Based Estimation

The results in the previous sections show that in both Germany and Sweden, men's earnings increase more than women's when couples move across commuting zones. Additionally, we find that job layoffs increase the probability that a household moves across commuting zones by a greater degree when the man in the couple is laid off compared to the woman. There are numerous potential explanations for these results, but we focus on distinguishing between two of them: (1) men's higher potential earnings and greater returns to migration compared to women, and (2) a gender norm that prioritizes men's career advancement.

To distinguish between these two explanations, we develop a model of the household migration decision that extends the standard unitary household model by allowing the household to potentially place more weight on income earned by the man relative to the woman (Foged, 2016). We use the model to derive additional new empirical tests for whether or not the results in the previous sections can be rationalized with a standard unitary model with

gender differences in potential earnings¹³.

After presenting our theoretical results, we report additional empirical results that are directly motivated by the model, and we estimate the model parameters – separately for each country – using these additional empirical results and other moments from the data. We then use the model parameters to test (and reject) the unitary model in both countries, finding larger deviations in Germany as compared to Sweden. Lastly, we use the estimated model parameters to simulate the effects of job layoffs on migration (to compare to the estimated effects of job layoffs documented above), as well as the earnings effects of childbirth (the so-called “child penalty”).

2.5.1 Model

Model setup. There is a unit mass of households, each household has a male ($i = M$) and a female ($i = F$), and there are two periods ($t = 1, 2$). Households decide whether or not to move between the two periods. Income in period 1 represents each individual’s pre-move permanent income and is assumed to be drawn independently from a log-normal income distribution: $\log(y_{i1}) \sim N(\mu_i, \sigma^2)$.¹⁴ With this setup, there is an average gender gap in period 1 of $\exp(\mu_M + \sigma^2/2) - \exp(\mu_F + \sigma^2/2)$. Define $s = y_{F1}/(y_{M1} + y_{F1})$ to be the female’s share of total household income in period 1.

Migration decision. For simplicity, we assume that each household member receives the same income in period 2 if the household chooses not to move. Each household member independently draws a potential income in period 2

¹³Like our model, Foged (2016) develops a model where households discount income earned by the wife relative to the husband, but the paper focuses on developing predictions about how the probability of moving varies with the female earnings share of household income, while we focus on how the expected change in income after moving varies with the female earnings share. As we show in the Appendix using simulations, the predictions in Foged (2016) on how the probability of moving varies with the female earnings share is sensitive to functional form assumptions and is not robust to extensions for assortative mating, while our simulations show that our predictions in Propositions 1 and 2 are robust to both of these extensions. As a result, we conclude that the earnings effects of migration are a more robust and reliable way to infer whether or not households discount income earned by the wife relative to the husband.

¹⁴This baseline setup implicitly assumes no assortative mating and assumes that the log income distributions for men and women have equal variances. We relax both of these assumptions in the Appendix and show in simulations that our main propositions go through with both of these extensions.

that they would receive if they move, with $y_{i2} = (1 + \varepsilon_{i2})y_{i1}$ and $\varepsilon_{i2} \sim N(\mu_r, \sigma_r^2)$. The μ_r and σ_r parameters capture heterogeneity in the returns to migration, and we assume that the average return to moving is the same across genders when expressed as a percentage of baseline income. We assume that a **unitary household** chooses to move if and only if the increase in household income from moving is greater than the household's (money-metric) utility cost of moving c . Define the change in income for each household member as $\Delta y_i = y_{i2} - y_{i1}$. With this setup, a unitary household moves if $\Delta y_M + \Delta y_F > c$. A **non-unitary household** places relatively less weight on the female's income by a share parameter β (with $0 < \beta < 1$); this type of household will move if $\Delta y_M + \beta \Delta y_F > c$.

The following proposition describes the expected return to moving (conditional on moving) in the full population:

Proposition 1 *If $\mu_M > \mu_F$ and all households are unitary households, then the expected return to moving (conditional on moving) is larger for men than women: $E[\Delta y_M - \Delta y_F | \Delta y_M + \Delta y_F > c] > 0$.*

Proof. See Appendix.

This proposition shows that if there is a baseline gender gap and the returns to migration are (assumed to be) the same for both genders, then this implies that in unitary households men will systematically benefit from moving relative to women.

Intuitively, it is more likely that the male household member draws a potential income in period 2 that exceeds the household's cost of moving, and so conditional on moving, it is more likely that the move is a move that benefits the man rather than the woman. This implies that the previous reduced-form empirical results on their own do not reject a standard unitary model and do not necessarily imply any inefficiency in household decision-making.

The full proof is given in the Appendix, but some intuition can be gained from the following lemma:

Lemma 1 *If $\mu_M > \mu_F$ and all households are unitary households, then the expected return to moving (conditional on moving) is larger for men than women for any household with $0 < s < 0.5$; i.e., for all $0 < s < 0.5$, $E[\Delta y_M - \Delta y_F | s, \Delta y_M + \Delta y_F > c] > 0$.*

Proof. See Appendix.

Lemma 1 says that for any household with $0 < s < 0.5$, the expected return to moving is larger for men than women. Since $\mu_M > \mu_F$ and there is no assortative mating in our baseline model, then $E[s] < 0.5$ in the population. As a result, integrating across all households in the population ends up with an unconditional average return that is larger for men than women.

While Proposition 1 shows that it is not possible to rule out a unitary model based on the gender gap in expected returns to migration (among the households who choose to move), the next proposition shows that for the households at $s = 0.5$, the expected return to moving (conditional on moving) is the same for men and women:

Proposition 2 *If $\mu_M > \mu_F$ and all households are unitary households, then the expected return to moving (conditional on moving) for men and women is equal for households at $s = 0.5$; i.e., $E[\Delta y_M - \Delta y_F | s = 0.5, \Delta y_M + \Delta y_F > c] = 0$.*

Proof. *See Appendix.*

Proposition 2 shows that our model with unitary households makes a sharp prediction for households at $s = 0.5$. For these households, when two spouses have identical income in period 1 and the same distribution of potential returns to moving, the result is that it is equally likely that each member ends up being the “trailing spouse” when the household chooses to move. Intuitively, for the couples with $s = 0.5$, the probability of drawing a potential income that exceeds the household’s mobility cost is the same for each household member. It is therefore equally likely that a move benefits the man as it benefits the woman.

Propositions 1 and 2 are both established in a very simplified setting, with baseline log income distributions for men and women having equal variance (homoskedasticity), and no assortative mating. The Appendix presents proofs and simulations of extended versions of the baseline model that allow for unequal variances across genders in baseline log income and also allow for assortative mating, and both results carry through with these model extensions.

We now turn to non-unitary households, where households behave “as if” they put less weight on income earned by the woman relative to income earned by the man, and this relative weight is given by the parameter β , with $0 < \beta < 1$ (so that $\beta = 1$ corresponds to a standard unitary household).

In contrast to Proposition 2, when households are non-unitary households with $0 < \beta < 1$, the expected return to moving (conditional on moving) is larger for men compared to women at $s = 0.5$, with the gap decreasing as β approaches 1.

Proposition 3 *If $\mu_M > \mu_F$ and all households are non-unitary households with $0 < \beta < 1$, then the expected return to moving (conditional on moving) is larger for men than women for households at $s = 0.5$; i.e., $E[\Delta y_M - \Delta y_F | s = 0.5, \Delta y_M + \beta \Delta y_F > c] > 0$, with the expectation approaching 0 as β approaches 1 from below.*

Proof. *See Appendix.*

Proposition 3 shows that an empirical implication of the unitary household model is that we should be able to find households with similar income and potential returns from moving, and these households should on average have returns to moving (conditional on moving) that are similar by gender. If we continue to find (within the set of households at $s = 0.5$) that men disproportionately benefit from moving compared to women, then we will conclude that the household's behavior is not consistent with a unitary model and conclude instead that households put less weight on income earned by the woman, with $0 < \beta < 1$.

These propositions thus make clear that men disproportionately benefiting from migration does not on its own conflict with predictions from a standard unitary household model when there are pre-existing gender earnings gaps. Intuitively, if the returns to migration are similar across the income distribution (in percentage terms), then men and women who move as couples will tend to experience increased earnings inequality within the household. In order to rule out a unitary model, we need to “zoom in” on the households near $s = 0.5$.

These theoretical results therefore motivate additional empirical specifications testing for heterogeneity in the effects of migration by the female share of household income prior to the move. Specifically, they imply we should expand the earnings regression models that estimate the earnings effects of migration to estimate how the earnings effects of migration vary with s .

2.5.2 Heterogeneity in the Effects of Migration on Earnings by Female Share of Household Income

Our results based on the full sample indicate that men realize significant positive returns from moving, while women are more likely to leave the workforce in the first years after the move. Based on the results in the previous subsection, we now examine how the returns to moving differ based on each individual's predicted share of household income.

In order to operationalize the additional empirical tests suggested by the model, we first construct a measure of (predicted) female share of household income. To do this, we estimate predicted income from a regression model. Specifically, we run a regression on a random sample of the full population of employed individuals in each country aged 25-54. The regression model relates log annual earnings to a large set of controls: potential experience dummies, child dummy, education dummies, and year dummies.¹⁵ In Sweden, we also include detailed indicators for the college majors for the individuals who attended either college and vocational training, and we interact these college major indicators with the education dummies in the prediction model. In Germany, we use first occupation instead of college majors.

We then use these regression models to construct a measure of predicted income in the year prior to the move for each member of the household, and we calculate the predicted female share of household income in both of our samples. Figures 2.B.6 and 2.B.7 show the distribution of predicted incomes for the men and women in our sample, and the predicted female share using this prediction model. We use the predicted female share of household income (\hat{s}) as our empirical proxy for the s in the model.

We choose to use predicted female share rather than the actual share partly because our layoff results indicate a clear gender-specific effect of layoffs on the probability of moving, so women with very high income shares in the years right before a move may be disproportionately made up of households where the man was recently laid off. In these households, the fact that the man disproportionately benefits from moving could mechanically come from a kind of "mean revision" arising from the layoff event that occurred prior to the migration decision. Additionally, actual earnings may not reflect

¹⁵The three education levels we use are high school, vocational training, and college.

an individual's true earnings potential, particularly for women; for example, Bertrand et al. (2015) find that relative income concerns affect actual earnings, as women may prefer to earn less to avoid out-earning their spouses. Our use of a predicted female earnings share measure is designed to address both of these concerns.¹⁶



As a result, we focus on households with similar predicted income based on education and experience, and we assume that households with similar potential income have similar returns to migration. These are the households we want to “zoom in” on in order to estimate the earnings effects for households at or near $\hat{s} = 0.5$.

To get an initial sense of how the earnings effects of moving vary with \hat{s} , we first divide our sample into couples where the man has the higher predicted share of household income and those where the woman has the higher share. The results are shown in Figure 2.6. This figure shows that the gap between men's and women's wage income is a bit smaller when women have a larger (i.e., greater than 50%) predicted share of household income, although the point estimates suggest that women still earn slightly less than men. This implies that on average men benefit more from relocation than women for households with $\hat{s} < 0.5$, but women do not benefit more from relocation than men for households with $\hat{s} > 0.5$. Appendix Figure 2.B.4 shows that we still do not find evidence that women benefit more from relocation than men when $\hat{s} > 0.5$ using a gender-specific measure of predicted income (as compared to the gender-blind prediction that we use in our baseline analysis). Taken together, these results are our first pieces of evidence that $\beta < 1$ in both countries.



¹⁶Additionally, in our model where households behave “as if” they value the income earned by the woman less than the income earned by the man, women may choose to work less and earn less precisely because of this “discounting” of the woman's income within the household. That is, even when men and women have the same potential income, there will be a gender earnings gap within the household when $\beta < 1$ allowing for endogenous labor supply responses.

Figure 2.6: Impact of Move on Wage Income – By Gender-blind Predicted Female Share of HH Income

(a) Female Share of HH Income < 50%, Ger- (b) Female Share of HH Income < 50%, Swe-
many den



(c) Female Share of HH Income \geq 50%, Ger- (d) Female Share of HH Income \geq 50%, Swe-
many den



Notes: This figure displays the event study results that estimate the effect of moving on wage income in each year relative to the year before the move ($t - 1$). Each point estimate has a corresponding 95% confidence interval calculated using standard errors clustered at the individual level. The regressions are run separately by gender. The coefficients and standard errors (in parentheses) in the upper right corner of each figure are 6 and 11-year averages of the post-move point estimates (from $t = 0$ to $t = 5$ and $t = 10$), in this order, for men (M) and women (W). Gender-blind predicted earnings are calculated by regressing men's log individual income on experience indicators and education level interacted with field of study, in a way that men and women with the same covariates have the same predicted wage income.

In order to estimate how the earnings effects of migration vary with \hat{s} , we estimate flexible spline specifications that interact spline functions of \hat{s} with indicator variables capturing the years after the move. The spline specifications are used to construct predicted values of the average earnings effects of migration at $\hat{s} = 0.4$ and $\hat{s} = 0.5$ for men and women. These results are summarized in Table 2.4 below. Columns (1) and (2) show the results for Germany, and columns (3) and (4) show the results for Sweden.

Table 2.4: How Do the Effects of Moving by Gender Vary with the Predicted Female Share of Household Income?

Predicted Female Share of Household Income, \hat{s}	Germany		Sweden	
	Men (1)	Women (2)	Men (3)	Women (4)
$\hat{s} = 0.4$	7.95 (.)	0.90 (.)	3.18 (.)	-0.14 (.)
$\hat{s} = 0.5$	5.84 (.)	2.20 (.)	1.01 (.)	0.675 (.)
Full sample	8.24 (0.35)	1.91 (0.18)	3.29 (0.12)	1.24 (0.08)

Notes: This table presents estimates from spline regressions on the earnings effects of moving by gender, allowing for the effects of moving to vary with the predicted female share of household income. The final row reports the estimates from the full sample for comparison. These results are based on gender-blind earnings predictions.

Comparing across the columns, we see that in Sweden there is a smaller baseline gender gap (in the years prior to migration) and a lower migration rate compared to Germany. In both countries, at $\hat{s} = 0.4$ there are large differences by gender. For these households, men's earnings increase by 10-15 percent in both countries, while women's income actually declines in Sweden and does not change in Germany.

Note the model described above can generate average declines in earnings for women even in a unitary model if there is a large variance in idiosyncratic mobility costs across households. Intuitively, if there are many other reasons why households move besides to increase labor earnings, then sometimes one or both members in the household will choose to move even though their income declines, and this is more likely to happen for women compared to

men at $\hat{s} = 0.4$, even in unitary households.¹⁷

Turning to $\hat{s} = 0.5$, we see in both countries the average return to migration is lower for men and higher for women (compared to $\hat{s} = 0.4$), but a gender gap remains at $\hat{s} = 0.5$ in both countries. This is another piece of evidence against a standard unitary model explaining our results. The unitary model would predict at $\hat{s} = 0.5$ that the average return to migration should be the same for men and women. The gap does “converge more” between men and women in Sweden compared to Germany, which is our first piece of evidence that households may deviate less from unitary model in Sweden compared to Germany (i.e., β is closer to 1).

2.5.3 Model-Based Estimation

We now use the moments and estimates in the table to estimate the model parameters. We first calibrate the baseline distribution of income prior to migration in both countries. This requires fitting log normal income distribution for men and women in both countries. These results are reported in Panel A of Table 2.5. Consistent with the results in Table 2.4 in the previous subsection, there is a larger baseline gender gap in Germany as compared to Sweden.

With these parameters calibrated, there are five remaining model parameters: the mean and variance parameters governing the returns to migration for men and women (μ_m and σ_m), the mean and variance parameters governing the household’s idiosyncratic mobility cost (μ_c and σ_c), and the non-unitary household parameter β .

To identify and estimate these five model parameters, we use the following five moments: the migration rate (share of households moving during the sample period), the average returns to migration for men and women at $\hat{s} = 0.4$, and the average returns to migration for men and women at $\hat{s} = 0.5$. Intuitively, if $\beta = 1$, then the average returns to migration for men and women at $\hat{s} = 0.5$ should be the same, so we lose one moment and one parameter. This tells us that the “gap” in average returns to migration at $\hat{s} = 0.5$ primarily identifies the parameter β . Varying σ_c changes the average migration rate,

¹⁷While the theoretical model focuses on a single mobility cost parameter c , in our model-based estimation we will allow for heterogeneity in household mobility costs by specifying a distribution of mobility costs alongside a distribution of returns to migration.

Table 2.5: Model Parameter Estimates

	Germany	Sweden
	(1)	(2)
Panel A: Baseline log normal income distribution parameters		
Mean log income, men	3.78	2.99
Standard deviation of log income, men	0.43	0.51
Mean log income, women	3.37	2.81
Standard deviation of log income, women	0.67	0.70
Panel B: Estimated model parameters		
Mean returns to migration, μ_r	-0.10	-0.32
Standard deviation in the returns to migration, σ_r	0.34	0.25
Mean household mobility cost, μ_c	0.25	-0.15
Standard deviation of household mobility cost, σ_c	0.15	1.43
Relative weight on woman's income compared to man's income, β	0.63	0.86

Notes: Panel A displays the mean and standard deviation of log income in the year prior to the move for the full sample of movers. These values are used to calibrate the parameters of the log normal income distribution. Panel B displays the model-based estimates for both countries based on a simple equal-weighted minimum distance estimator, using as moments the average migration rate and the effects of moving at $\hat{s} = 0.4$ and $\hat{s} = 0.5$ reported in Table 2.4.

but does not affect the average returns to migration (conditional on moving), so the migration rate primarily identifies the parameter σ_c . The identification of the other three parameters is more subtle, but they are jointly identified by the relative gaps between men and women at $\hat{s} = 0.4$ compared to $\hat{s} = 0.5$, given β .

To estimate the model parameters, we simulate the model a large number of times and search for the combination of model parameters that minimize the sum of the squared distance between the moments and the simulated values of the moments from the model (since σ_c can always be chosen to target a given migration rate, we search over the other four parameters and then choose σ_c to match migration rate exactly). The model-based parameters are reported in Panel B of Table 2.5. Table 2.6 compares the actual moments and the simulated moments at the chosen model parameters, which shows that the model has a good fit in both countries.

Turning to the estimated parameters, we see that the returns to migration is “shifted down” in Sweden compared to Germany, which is consistent with both the lower estimated average returns to moving. Since the migration rate is not that much lower in Sweden, however, we need the mobility costs to be more heterogeneous in Sweden in order to generate enough migration in the model to match the data.

Our primary parameter of interest is the β parameter, which is estimated to be $\beta = 0.86$ in Sweden and $\beta = 0.56$. One way to assess the importance of $\beta < 1$ is to re-simulate the model with $\beta = 1$, holding the other parameters constant. Panel C of Table 2.6 shows that this results in a worse model fit, particularly for the $\hat{s} = 0.5$ households. An alternative is to re-estimate the model restricting $\beta = 1$; column (4) of Table 2.6 shows that this model also has a worse fit, particularly for Germany.

The conclusion from the model-based estimation is therefore that the earnings effects of migration in both countries are difficult to reconcile with a standard unitary household model, and the earnings effects at different predicted female shares of household income suggest that households in both countries place less weight on income earned by woman compared to man, particularly in Germany.

The larger departure from the unitary model in Germany is interesting because Germany also has a larger baseline gender gap (and, as we discuss

later, a larger female “child penalty”). This raises the possibility that the baseline gender gap itself may be due to the same factors that lead households to seemingly “under-react” to women’s potential returns from relocation. We conclude this section by using the estimated model to carry out two additional exercises: we use the model to simulate the effects of job layoffs on migration and the effects of childbirth on earnings.

Table 2.6: Assessing Model Fit

Predicted Female Share of Household Income, \hat{s}	Germany		Sweden	
	Men (1)	Women (2)	Men (3)	Women (4)
Panel A: Empirical Estimates				
$\hat{s} = 0.4$	7.95	0.90	3.18	-0.14
$\hat{s} = 0.5$	5.84	2.20	1.01	0.68
Panel B: Simulated Moments from Baseline Model				
$\hat{s} = 0.4$	7.47	1.02	3.54	-0.40
$\hat{s} = 0.5$	5.23	2.33	1.25	0.74
χ^2 goodness-of-fit statistic	0.116		0.587	
Panel C: Simulated Moments Setting $\beta = 1$ (holding other parameters constant)				
$\hat{s} = 0.4$	6.58	2.34	3.33	-0.93
$\hat{s} = 0.5$	3.96	3.70	0.94	1.22
χ^2 goodness-of-fit statistic	4.168		4.899	
Panel D: Simulated Moments Restricting to $\beta = 1$ (re-estimating other parameters)				
$\hat{s} = 0.4$	7.34	1.53	2.81	-0.62
$\hat{s} = 0.5$	4.17	4.07	1.15	0.71
χ^2 goodness-of-fit statistic	2.564		1.703	

Notes: This table presents the empirical estimates of the effects of moving at $\hat{s} = 0.4$ and $\hat{s} = 0.5$ and compares to the baseline model estimates and alternative model estimates setting $\beta = 1$ and either holding other parameters constant or re-estimating the other model parameters.

2.5.4 Additional Implications of $\beta < 1$: Gender Differences in Effects of Job Layoffs on Relocation and Gender Differences in Child Penalties

An additional way to assess the fit of the model with the estimated $\beta < 1$ parameter is to simulate an exogenous decline in male or female income (from job separation caused by mass layoff), and then predict the change in the probability of moving depending on whether or not the male or female was laid off. We can then compare these results to the reduced-form effects above. This is an “out-of-sample” test of model fit because the effects of job layoff on the probability of relocating by gender were not directly targeted in the model-based estimation. To do this, we simulate the model at the parameters estimated in each country and we exogenously reduce income by the man or woman by 20 percent and then estimate the resulting change in the probability of moving. The results in Panel A of Table 2.7 show that the model can accurately reproduce a gender gap in the effects of a job layoff on the probability of moving. The model somewhat under-predicts the gender gap in Germany and somewhat over-predicts the gender gap in Sweden, but this could come from the fact that we currently assume the exogenous income change is the same in both countries.

Lastly, we use our estimated model to simulate the change in earnings following the birth of the couple’s first child to see how much our estimated $\beta < 1$ parameter can account for the female “child penalty” in both countries. Specifically, we compare our simulated results to the results from Kleven et al. (2019a) that estimate the child penalty in a large number of countries. They find that the child penalty is much larger in Germany, and we also find a larger departure from $\beta = 1$ based on the earnings responses to relocation. One interpretation of the child penalty is that the household puts less weight on income declines by the woman (as compared to the man), which means that even if the man and woman in a household have equal ability in child-rearing, the household may still choose to have the woman reduce her hours in the formal labor market. The Appendix formalizes this argument and shows that the child penalty should be closely related to $(1 - \beta)$ in this case. We do this by extending the model in Andresen and Nix (2022) to allow for $\beta < 1$, and Panel B of Table 2.7 shows that this simulated model can account for most

of the female “child penalty” in both Germany and Sweden, and can also account for most of the difference between Germany and Sweden. In other words, the greater deviation from unitary model in Germany can account for most of the larger child penalty according to our simulated model.

Table 2.7: Model-Based Simulations

	Germany		Sweden	
	Men (1)	Women (2)	Men (3)	Women (4)
Panel A: Proportional Change in Probability of Moving After Layoff				
Empirical estimate	1.83	1.24	1.89	0.88
Model-based simulation	1.61	1.29	1.95	1.58
Panel B: Proportional Change in Earnings After Birth of First Child				
Empirical estimate from Kleven et al. (2019b)	-0.02	-0.61	-0.06	-0.26
Model-based simulation	-0.04	-0.48	-0.07	-0.19
Implied share of Female “child penalty” accounted for the country-specific β estimate		78.7%		71.5%

Notes: Panel A uses baseline model-based estimates to simulate changes in the probability of moving after an exogenous job displacement. Panel B simulates change in earnings after birth of first of child to compare the implied changes (at estimated country-specific β) to the actual changes estimated in Kleven et al. (2019b).

2.6 Additional Evidence and Alternative Explanations

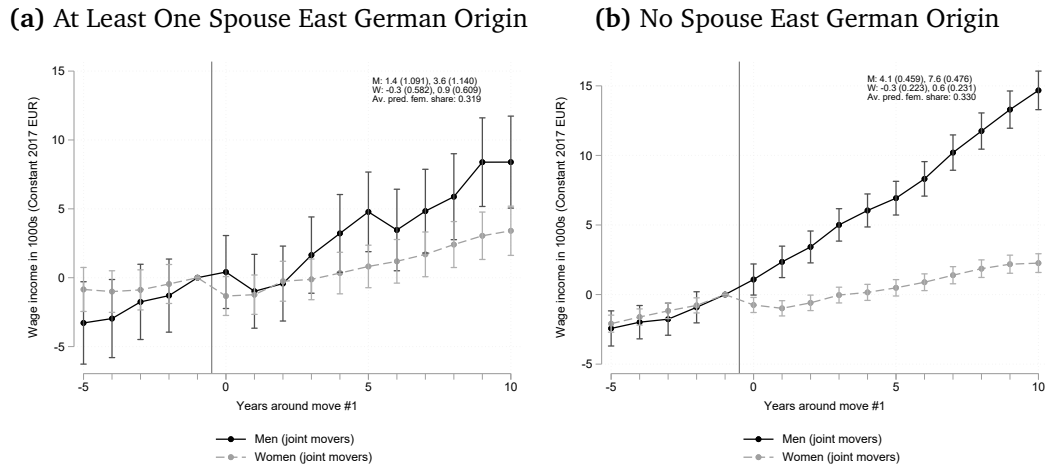
The previous section argues that the earnings gap cannot be rationalized by a unitary household model in which couples maximize household income, and is better explained by a model in which couples adhere to a gender norm that essentially means putting less weight on women’s earnings. In this section we provide additional evidence that the results are driven by gender norms, and rule out several alternative interpretations.

2.6.1 Gender Norms: East and West Germany

As a more direct test of whether culture or norms explain part of the earnings gap, we test whether the gap varies based on whether men and women have origins in East or West Germany. East Germany has had high rates of female labor force participation due to its history as a socialist state where women were strongly encouraged to work. Existing research has shown that whether women grow up in East or West Germany influences decisions concerning labor supply (Boelmann et al., 2021). We use variation in couples' family origins as a source of variation in gender norms.

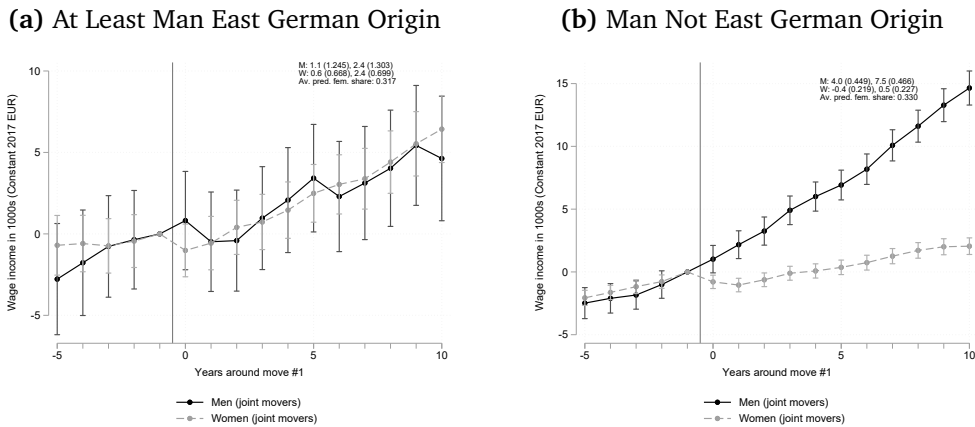
Figure 2.7 and 2.8 show that especially for more egalitarian couples returns to moving are more equal for men and women.

Figure 2.7: East vs. West German Origin



Notes: This figure displays the event study results that estimate the effect of moving on different outcomes in each year relative to the year before the move ($t - 1$) for different samples depending on spouses' origin. Each point estimate has a corresponding 95% confidence interval calculated using standard errors clustered at the individual level. The regressions are run separately by gender. The coefficients and standard errors (in parentheses) in the upper right corner of each figure are 6 and 11-year averages of the post-move point estimates (from $t = 0$ to $t = 5$ and $t = 10$), in this order, for men (M) and women (W).

Figure 2.8: East vs. West German Origin, Men



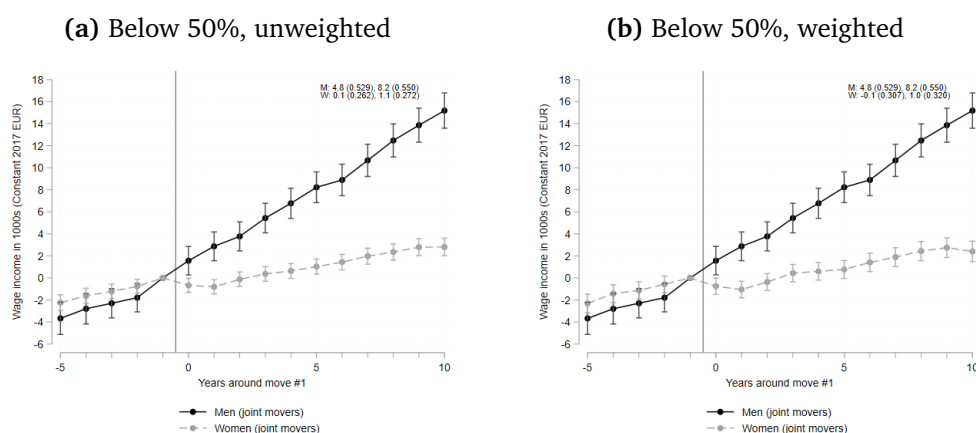
Notes: This figure displays the event study results that estimate the effect of moving on different outcomes in each year relative to the year before the move ($t - 1$) for different samples depending on man's origin. Each point estimate has a corresponding 95% confidence interval calculated using standard errors clustered at the individual level. The regressions are run separately by gender. The coefficients and standard errors (in parentheses) in the upper right corner of each figure are 6 and 11-year averages of the post-move point estimates (from $t = 0$ to $t = 5$ and $t = 10$), in this order, for men (M) and women (W).

2.6.2 Alternative Explanations

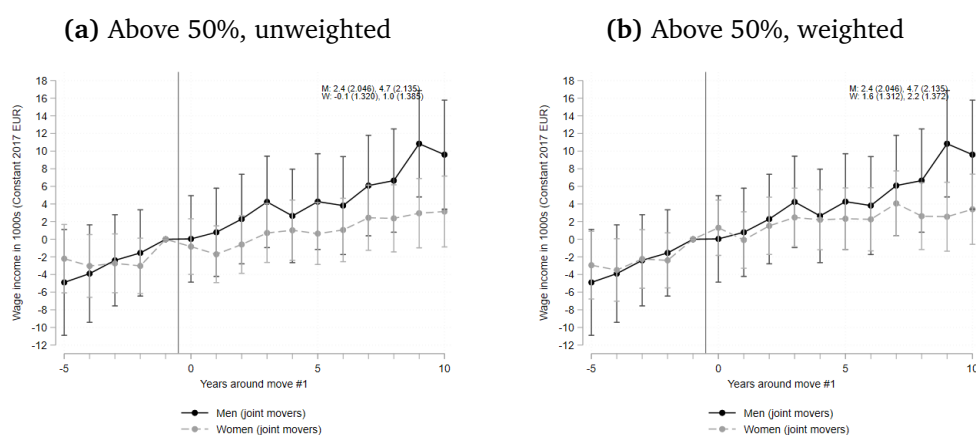
We now explore several alternative explanations for our findings. First, we test whether the results are driven by women selecting into occupations that have lower returns to moving. Second, we explore the possibility that women's lower returns to moving are made up for by a non-wage amenity.

Given that women tend to be in occupations with lower wage growth, it is possible that these same jobs have lower returns to moving. To test whether this explains the results, we estimate our event study equation but reweight the sample so that women have the same occupation distribution as men.¹⁸ Figure 2.9 panel (a) shows the results for couples in which women earn less than 50% of household income and figure 2.10 panel (a) for couples in which women earn more than 50% of household income. We also include the unweighted regression results for comparison (panels (b) respectively). The results are largely unchanged, suggesting that occupational sorting and differences in returns to moving for these occupations are not driving the results.

¹⁸To do so, we limit our movers sample to couples in which both individuals are working in occupations with at least 10 individuals in the occupation, again within our movers sample. We further restrict to occupations that have at least one man and one woman. We then re-weight so that the women in the sample have the same occupation distribution as men. Occupations are defined at the 4-digit level for this analysis.

Figure 2.9: Wage Income Results by Occupation - Below 50%, Germany


Notes: This figure displays the event study results that estimate the effect of moving on different outcomes in each year relative to the year before the move ($t - 1$) accounting for differences in occupations between men and women. Each point estimate has a corresponding 95% confidence interval calculated using standard errors clustered at the individual level. The regressions are run separately by gender. The coefficients and standard errors (in parentheses) in the upper right corner of each figure are 6 and 11-year averages of the post-move point estimates (from $t = 0$ to $t = 5$ and $t = 10$), in this order, for men (M) and women (W).

Figure 2.10: Wage Income Results by Occupation - Above 50%, Germany

Notes: This figure displays the event study results that estimate the effect of moving on different outcomes in each year relative to the year before the move ($t - 1$) accounting for differences in occupations between men and women. Each point estimate has a corresponding 95% confidence interval calculated using standard errors clustered at the individual level. The regressions are run separately by gender. The coefficients and standard errors (in parentheses) in the upper right corner of each figure are 6 and 11-year averages of the post-move point estimates (from $t = 0$ to $t = 5$ and $t = 10$), in this order, for men (M) and women (W).

It is also possible that women's returns to moving come in the form of non-wage amenities. For example, research has shown that women choose jobs with shorter commute times (Le Barbanchon et al., 2020). A couple could therefore be treating each member equally but women benefit from a shorter commute whereas men benefit from a higher salary or wage. To explore this possibility, we look at how distance to work changes following a move. Figure 2.11 shows that, while men's average commute increases slightly, women's average distance from work does not change. It is possible that women are moving to firms that are offering other non-wage amenities, but we are unable to test for this in our data.

Figure 2.11: Distance as outcome (in km) - Sweden



Notes: This figure displays means of commuting distances for men and women (in km) in Sweden.

2.7 Conclusion

Over the past half a century, women have made great strides in the labor market. However, despite substantial gender convergence, there are still large differences between men and women. In this paper, we investigate an aspect that contributes to gender differences in the labor market which has not received much attention in the recent literature: gender differences

in the returns to moving. Using administrative data from Germany and Sweden, we use an event study design to estimate the labor market effects of couples' long-distance moves, and we find that men's earnings increase significantly after a long-distance move, and women's earnings increase by less (if at all). These results echo some of the results in previous studies (see, e.g., Blackburn, 2010a; Cooke et al., 2009; LeClere and McLaughlin, 1997; Sandell, 1977; Blackburn, 2010b; Cooke, 2003; Spitze, 1984; Rabe, 2009), but the unusually large and representative sample of opposite-sex couples in our analysis provides new evidence of this gender divergence. While we find that men benefit almost exclusively through higher wages, women's losses are mostly due to exiting the labor market or being employed for fewer days of the year.

Using a model of household decision-making where households "discount" the income earned by the woman compared to the man, we test and reject the unitary model in both countries, with larger departures in Germany compared to Sweden. Overall, we conclude that a gender norm that prioritizes men's career advancement can simultaneously (and parsimoniously) account for three different gender differences in labor market outcomes: the earnings effects of relocation, the probability of moving following a job layoff, and the earnings effects of the birth of a child (the so-called "child penalty"). Of course, it is hard to fully rule out explanations based on gender differences in preferences (e.g., preferences for child-rearing, preferences for leisure, preferences for part-time work or flexible hours), but we interpret our model-based estimates as potentially suggesting a unifying explanation that households systematically pass up opportunities to maximize lifetime household income because households behave "as if" income earned by the woman is worth less than income earned by the man. If true, this is hard to square with many models of efficient household decision-making.

We conclude by briefly mentioning several areas of future work. First, we make several simplifying assumptions in the model. For example, we assume away heterogeneity in the β parameter. This is done to make the identification as transparent as possible, but it may be possible to estimate a richer model where β can vary with observed and unobserved household characteristics. Second, we focus on two countries with readily-available administrative data and fairly different labor market institutions, but we think our framework

can be easily implemented in other countries. If we are right that the female “child penalty” is driven at least in part by our β parameter, then one should see larger departures from the unitary model in countries with larger child penalties. Lastly, we conjecture that our model may be consistent with certain household bargaining models with limited commitment, and it would be interesting to try to make this connection more precise. For the questions addressed in this paper, we did not need a micro-foundation of where the $\beta < 1$ parameter is coming from, but for other questions it may be useful to give more details of exactly how the households come to treat women’s income as less valuable than men’s.

Bibliography

- Andresen, M. E. and Nix, E. (2022). What causes the child penalty? evidence from adopting and same-sex couples. *Journal of Labor economics*, 40(4):971–1004.
- Angelov, N., Johansson, P., and Lindahl, E. (2016). Parenthood and the gender gap in pay. *Journal of Labor Economics*, 34(3):545–579.
- Bartel, A. P. (1979). The migration decision: What role does job mobility play? *American Economic Review*, 69(5):775–86.
- Bertrand, M., Kamenica, E., and Pan, J. (2015). Gender Identity and Relative Income within Households *. *The Quarterly Journal of Economics*, 130(2):571–614.
- Blackburn, M. L. (2010a). The impact of internal migration on married couples' earnings in Britain. *Economica*, 77(307):584–603.
- Blackburn, M. L. (2010b). Internal migration and the earnings of married couples in the United States. *Journal of Economic Geography*, 10(1):87–111.
- Boelmann, B., Raute, A., and Schonberg, U. (2021). Wind of change? cultural determinants of maternal labor supply.
- Bolotnyy, V. and Emanuel, N. (2022). Why do women earn less than men? evidence from bus and train operators. *Journal of Labor Economics*.
- Burke, J. and Miller, A. R. (2018). The effects of job relocation on spousal careers: Evidence from military change of station moves. *Economic Inquiry*, 56(2):1261–1277.
- Bursztyjn, L., Fujiwara, T., and Pallais, A. (2017). 'acting wife': Marriage market incentives and labor market investments. *The American Economic Review*, 107(11):3288–3319.
- Buser, T., Niederle, M., and Oosterbeek, H. (2014). Gender, Competitiveness, and Career Choices. *The Quarterly Journal of Economics*, 129(3):1409–1447.

- Compton, J. and Pollak, R. A. (2007). Why Are Power Couples Increasingly Concentrated in Large Metropolitan Areas? *Journal of Labor Economics*, 25:475–512.
- Cooke, T. J. (2003). Family migration and the relative earnings of husbands and wives. *Annals of the Association of American Geographers*, 93(2):338–349.
- Cooke, T. J., Boyle, P., Couch, K., and Feijten, P. (2009). A longitudinal analysis of family migration and the gender gap in earnings in the united states and great britain. *Demography*, 46:147–167.
- Cortes, P. and Pan, J. (2022). Children and the remaining gender gaps in the labor market. *Journal of Economics Literature*, forthcoming.
- Cortes, P., Pan, J., Pilossoph, L., and Zafar, B. (2021). Gender differences in job search and the earnings gap: Evidence from business majors. *National Bureau of Economic Research Working Paper*, (28820).
- Costa, D. L. and Kahn, M. E. (2000). Power Couples: Changes in the Locational Choice of the College Educated, 1940–1990*. *The Quarterly Journal of Economics*, 115(4):1287–1315.
- Fadlon, I., Lyngse, F. P., and Nielsen, T. H. (2022). Causal effects of early career sorting on labor and marriage market choices: A foundation for gender disparities and norms. *NBER Working Paper 28245*.
- Foged, M. (2016). Family migration and relative earnings potentials. *Labour Economics*, 42:87–100.
- Frank, R. H. (1978). Why women earn less: The theory and estimation of differential overqualification. *The American Economic Review*, 68(3):360–373.
- Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review*, 104(4):1091–1119.
- Goldschmidt, D., Klosterhuber, W., and Schmieder, J. F. (2017). Identifying couples in Administrative Data. *Journal for Labour Market Research*, 50:29–43.

- Illing, H., Schmieder, J. F., and Trenkle, S. (2021). The gender gap in earnings losses after job displacement. Technical report, National Bureau of Economic Research.
- Kleven, H., Landais, C., Posch, J., Steinhauer, A., and Zweimueller, J. (2019a). Child penalties across countries: Evidence and explanations. *AEA Papers and Proceedings*, 109:122–126.
- Kleven, H., Landais, C., and SØgaard, J. E. (2019b). Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics*, 11(4):181–209.
- Kosfeld, R. and Werner, A. (2012). Deutsche arbeitsmarktregionen-neuabgrenzung nach den kreisgebietsreformen 2007–2011. *Raumforschung und Raumordnung— Spatial Research and Planning*, 70(1):49–64.
- Le Barbanchon, T., Rathelot, R., and Roulet, A. (2020). Gender Differences in Job Search: Trading off Commute against Wage*. *The Quarterly Journal of Economics*, 136(1):381–426.
- LeClere, F. B. and McLaughlin, D. K. (1997). Family Migration and Changes in Women’s Earnings: A Decomposition Analysis. *Population Research and Policy Review*, 16:315–335.
- Mincer, J. (1978). Family migration decisions. *Journal of Political Economy*, 86(5):749–773.
- Nivalainen, S. (2004). Determinants of Family Migration: Short Moves vs. Long Moves. *Journal of Population Economics*, 17:157–175.
- Polachek, S. and Horvath, F. (2012). A life cycle approach to migration: Analysis of the perspicacious peregrinator. *Research in Labor Economics*, 35:349–395.
- Rabe, B. (2009). Dual-earner migration. earnings gains, employment and self-selection. *Journal of Population Economics*, 24:477–497.
- Sandell, S. H. (1977). Women and the economics of family migration. *The Review of Economics and Statistics*, 59(4):406–414.

Spitze, G. (1984). The effect of family migration on wives' employment: How long does it last? *Social Science Quarterly*, 65(1):21.

Venator, J. (2020). Dual-earner migration decisions, earnings, and unemployment insurance. In *2020 APPAM Fall Research Conference*. APPAM.

Wiswall, M. and Zafar, B. (2017). Preference for the Workplace, Investment in Human Capital, and Gender. *The Quarterly Journal of Economics*, 133(1):457–507.

Appendix

2.A Proofs of Theoretical Results in Main Text

Proposition 4 *If $\mu_M > \mu_F$ and all households are unitary households, then the expected return to moving (conditional on moving) is larger for men than women: $E[\Delta y_M - \Delta y_F | \Delta y_M + \Delta y_F > c] > 0$.*

Proof. We want to show the following integral is positive, where $f(s)$ is the pdf of s :

$$\int_0^1 E[\Delta y_M - \Delta y_F | \Delta y_M + \Delta y_F > c] \cdot f(s) ds$$

Rewriting with the simplified form of the expression, we have:

$$\begin{aligned} \int_0^1 (1-2s) \left[\mu_r y_1 + \lambda \left(\frac{c - \mu_r y_1}{\sigma_r y_1 \sqrt{(1-s)^2 + s^2}} \right) \left[\frac{\sigma_r y_1}{\sqrt{(1-s)^2 + s^2}} \right] \right] \cdot f(s) ds = \\ = \underbrace{\int_0^1 (1-2s) \mu_r y_1 \cdot f(s) ds}_A + \underbrace{\int_0^1 (1-2s) \lambda \left(\frac{c - \mu_r y_1}{\sigma_r y_1 \sqrt{(1-s)^2 + s^2}} \right) \left[\frac{\sigma_r y_1}{\sqrt{(1-s)^2 + s^2}} \right] \cdot f(s) ds}_B \end{aligned}$$

We start with the first part of the expression, integral A. Assuming $s \in [0, 1]$, then $\int_0^1 f(s) ds = 1$.

$$\begin{aligned} \int_0^1 (1-2s) \mu_r y_1 \cdot f(s) ds &= \mu_r y_1 \int_0^1 (1-2s) f(s) ds \\ &= \mu_r y_1 \left[\int_0^{0.5} (1-2s) f(s) ds + \int_{0.5}^1 (1-2s) f(s) ds \right] \end{aligned}$$

We take the second integral from the expression above and integrate by substitution. Let $x = 1 - s$ and $dx = -ds$.

$$\begin{aligned} \int_{0.5}^1 (1-2s) f(s) ds &= \int_{0.5}^0 (1-2(1-x)) f(x) (-1) dx \\ &= \int_{0.5}^0 (-1)(1-2x) f(1-x) (-1) dx \\ &= - \int_0^{0.5} (1-2x) f(1-x) dx \end{aligned}$$

Returning to integral A:

$$\begin{aligned} \int_0^1 (1-2s)\mu_r y_1 \cdot f(s) ds &= \mu_r y_1 \left[\int_0^{0.5} (1-2s)f(s) ds + \int_{0.5}^1 (1-2s)f(s) ds \right] \\ &= \mu_r y_1 \left[\int_0^{0.5} (1-2s)f(s) ds - \int_0^{0.5} (1-2x)f(1-x) dx \right] \end{aligned}$$

We can combine the integrals in the last line because they have the same bounds of integration. Additionally, in the second integral, we defined the variable x , but the name of the variable itself is arbitrary so we can change it back to s for simplicity.¹⁹

$$\int_0^1 (1-2s)\mu_r y_1 \cdot f(s) ds = \mu_r y_1 \left[\int_0^{0.5} (1-2s)[f(s) - f(1-s)] ds \right]$$

Recall that if $f(x) \geq 0$ for $x \in [a, b]$, then $\int_a^b f(x) dx \geq 0$. In this case, we want to show that the function we are integrating is positive. Note that μ_r and y_1 are positive because they are the mean of the second period income and the first period household income, respectively. Additionally, $(1-2s)$ is positive between $(0, 0.5]$. Thus, for integral A to be positive, we have to show that $f(s) - f(1-s) > 0$.

The function, $f(s)$, is the PDF of s . To find the PDF of s , we have to determine its distribution. The first period incomes, y_{i1} for $i \in \{M, F\}$, have log-normal distributions, and s is a ratio of the incomes and has a logit-normal distribution, shown below.²⁰

$$\begin{aligned} s &= \frac{y_{F1}}{y_{F1} + y_{M1}} \\ &= \frac{1}{1 + y_{M1}/y_{F1}} \\ &= \frac{1}{1 + e^{\ln(y_{M1}/y_{F1})}} \\ &= \frac{1}{1 + e^{-[\ln(y_{F1}) - \ln(y_{M1})]}} \\ \implies f(s) &= \frac{1}{\sigma\sqrt{2\pi}} e^{-(\text{logit}(s) - \mu)^2 / (2\sigma^2)} \frac{1}{s(1-s)} \\ \mu &= \mu_F - \mu_M < 0 \\ \sigma &= 2\sigma^2 \end{aligned}$$

¹⁹Because we are considering s and x in separate integrals, we are able to do this. However, if s and x were within the same integral, and we were evaluating a double integral, then we would not be able to combined these integrals.

²⁰The logit-normal Pdf is defined only for $s \in (0, 1)$. Thus, to evaluate $f(s)$, we actually need to solve the improper integral between $(0, 1)$. Thus, for the rest of this proof, we will let $\int_0^1 f(s) ds = \int_{\rightarrow 0}^{\leftarrow 1} f(s) ds$. For our purposes, we will also assume that $f(0) = 0$ and $f(1) = 1$.

Plugging this back into integral A, we have:

$$\begin{aligned} f(s) - f(1-s) &= \frac{1}{\sigma\sqrt{2\pi}} \frac{1}{s(1-s)} \left[e^{-(\text{logit}(s)-\mu)^2/(2\sigma^2)} - e^{-(\text{logit}(1-s)-\mu)^2/(2\sigma^2)} \right] \\ &= \frac{1}{\sigma\sqrt{2\pi}} \frac{1}{s(1-s)} e^{-1/(2\sigma^2)} \left[e^{(\text{logit}(s)-\mu)^2} - e^{(\text{logit}(1-s)-\mu)^2} \right] \end{aligned}$$

To simplify the exponents of e , we use the following facts:

$$\begin{aligned} \text{logit}(s) &= \log\left(\frac{s}{1-s}\right) = \log(s) - \log(1-s) \\ \text{logit}(1-s) &= \log\left(\frac{1-s}{1-1+s}\right) = \log(1-s) - \log(s) \\ &= -\text{logit}(s) \end{aligned}$$

Let $\eta = \text{logit}(s)$. Returning to simplifying the expression for $f(s) - f(1-s)$:

$$\begin{aligned} f(s) - f(1-s) &= \frac{1}{\sigma\sqrt{2\pi}} \frac{1}{s(1-s)} e^{-1/(2\sigma^2)} \left[e^{\eta^2 - 2\mu\eta + \mu^2} - e^{(-\eta)^2 + 2\mu\eta + \mu^2} \right] \\ &= \frac{1}{\sigma\sqrt{2\pi}} \frac{1}{s(1-s)} e^{-1/(2\sigma^2) + \eta^2 + \mu^2} \left[e^{-2\mu\eta} - e^{2\mu\eta} \right] \\ \implies f(s) - f(1-s) &> 0 \end{aligned}$$

To summarize, considering all the components of integral A, we see that integral A is positive:

$$\int_0^1 (1-2s)\mu_r y_1 \cdot f(s) ds = \underbrace{\mu_r y_1}_{>0} \left[\int_0^{0.5} \underbrace{(1-2s)}_{>0} \underbrace{[f(s) - f(1-s)]}_{>0} ds \right] > 0$$

Now looking at integral B:

$$\int_0^1 (1-2s)\lambda \left(\frac{c - \mu_r y_1}{\sigma_r y_1 \sqrt{(1-s)^2 + s^2}} \right) \left[\frac{\sigma_r y_1}{\sqrt{(1-s)^2 + s^2}} \right] \cdot f(s) ds$$

Define $g(s) = \frac{k_1}{k_2 \sqrt{(1-s)^2 + s^2}}$ where k_1 and k_2 are constants. We want to show that the function C is symmetric over the line $x = 0.5$. This is equivalent to showing that $\int_0^{0.5} g(s) ds = \int_{0.5}^1 g(s) ds$.

$$\begin{aligned} \int_0^{0.5} \frac{k_1}{k_2 \sqrt{(1-s)^2 + s^2}} ds &= \frac{-k_1 \sinh^{-1}(1-2s)}{k_2 \sqrt{2}} \Big|_0^{0.5} = \frac{0.623225k_1}{k_2} \\ \int_{0.5}^1 \frac{k_1}{k_2 \sqrt{(1-s)^2 + s^2}} ds &= \frac{-k_1 \sinh^{-1}(1-2s)}{k_2 \sqrt{2}} \Big|_{0.5}^1 = \frac{0.623225k_1}{k_2} \end{aligned}$$

We can use this property of $g(s)$ to compare some of the terms in integral B. The terms,

$\lambda \left(\frac{c - \mu_r y_1}{\sigma_r y_1 \sqrt{(1-s)^2 + s^2}} \right)$ and $\frac{\sigma_r y_1}{\sqrt{(1-s)^2 + s^2}}$, can both be written in terms of $g(s)$ with different k_1 and k_2 . Given that $g(s)$ is symmetric about $x = 0.5$, we know that $\lambda \left(\frac{c - \mu_r y_1}{\sigma_r y_1 \sqrt{(1-s)^2 + s^2}} \right)$ and $\frac{\sigma_r y_1}{\sqrt{(1-s)^2 + s^2}}$ have the same values in the integrals when they are evaluated from $[0, 0.5]$ or $[0.5, 1]$.

Let $h(s) = \lambda \left(\frac{c - \mu_r y_1}{\sigma_r y_1 \sqrt{(1-s)^2 + s^2}} \right) \cdot \frac{\sigma_r y_1}{\sqrt{(1-s)^2 + s^2}}$. Then integral B can be rewritten as:

$$\int_0^1 (1-2s)h(s)f(s)ds = \int_0^{0.5} (1-2s)h(s)f(s)ds + \int_{0.5}^1 (1-2s)h(s)f(s)ds$$

Following the same steps for simplifying integral A, we integrate by substitution for the second integral above. Let $x = 1 - s$, $dx = -ds$.

$$\begin{aligned} \int_{0.5}^1 (1-2s)h(s)f(s)ds &= \int_{0.5}^0 (1-2(1-x))h(1-x)f(1-x)(-1)dx \\ &= - \int_0^{0.5} (1-2x)h(1-x)f(1-x)dx \end{aligned}$$

Combining the integrals:

$$\begin{aligned} \int_0^1 (1-2s)h(s)f(s)ds &= \int_0^{0.5} (1-2s)h(s)f(s)ds + \int_{0.5}^1 (1-2s)h(s)f(s)ds \\ &= \int_0^{0.5} (1-2s)h(s)f(s)ds - \int_0^{0.5} (1-2x)h(1-x)f(1-x)dx \\ &= \int_0^{0.5} (1-2s)[h(s)f(s) - h(1-s)f(1-s)]ds \end{aligned}$$

We have shown previously that $h(s)$ is symmetric about $s = 0.5$, so $h(s) = h(1-s)$. Therefore, whether integral B is positive depends on the sign of $f(s) - f(1-s)$. In simplifying integral A, we derived that $f(s) - f(1-s) > 0$, so this implies that integral B is also positive. Given that integral A and B are positive, this completes the proof that $\int_0^1 E[\Delta y_M - \Delta y_F | \Delta y_M + \Delta y_F > c] \cdot f(s)ds > 0$.

Lemma 2 *If $\mu_M > \mu_F$ and all households are unitary households, then the expected return to moving (conditional on moving) is larger for men than women for any household with $0 < s < 0.5$; i.e., for all $0 < s < 0.5$, $E[\Delta y_M - \Delta y_F | s, \Delta y_M + \Delta y_F > c] > 0$.*

Proof. To start, we expand the expectation, $E[\Delta y_M - \Delta y_F | s, \Delta y_M + \Delta y_F > c]$.

$$\begin{aligned}\Delta y_M - \Delta y_F &= (y_{M2} - y_{M1}) - (y_{F2} - y_{F1}) \\ &= (1 + \varepsilon_{M2})(1 - s)y_1 - (1 - s)y_1 - (1 + \varepsilon_{F2})sy_1 + sy_1 \\ &= \varepsilon_{M2}(1 - s)y_1 - \varepsilon_{F2}sy_1 \\ \Delta y_M + \Delta y_F &= (y_{M2} - y_{M1}) + (y_{F2} - y_{F1}) \\ &= (1 + \varepsilon_{M2})(1 - s)y_1 - (1 - s)y_1 + (1 + \varepsilon_{F2})sy_1 - sy_1 \\ &= \varepsilon_{M2}(1 - s)y_1 + \varepsilon_{F2}sy_1\end{aligned}$$

$$\implies E[\Delta y_M - \Delta y_F | s, \Delta y_M + \Delta y_F > c] = E[\varepsilon_{M2}(1 - s)y_1 - \varepsilon_{F2}sy_1 | s, \varepsilon_{M2}(1 - s)y_1 + \varepsilon_{F2}sy_1 > c]$$

We want to show that when $0 < s < 0.5$, $E[\varepsilon_{M2}(1 - s)y_1 - \varepsilon_{F2}sy_1 | \varepsilon_{M2}(1 - s)y_1 + \varepsilon_{F2}sy_1 > c] > 0$. Let $X = \varepsilon_{M2}(1 - s)y_1$ and $Y = \varepsilon_{F2}sy_1$, with their distributions defined below. Recall that $\varepsilon_{i2} \sim N(\mu_r, \sigma_r^2)$. We assume $cov(X, Y) = 0$.

$$\begin{aligned}X &= \varepsilon_{M2}(1 - s)y_1 & Y &= \varepsilon_{F2}sy_1 \\ &\sim N((1 - s)\mu_r y_1, ((1 - s)y_1 \sigma_r)^2) & &\sim N(s\mu_r y_1, (sy_1 \sigma_r)^2)\end{aligned}\quad (2.2)$$

With this substitution, we can rewrite the expectation to be $E[X - Y | X + Y > c]$, which allows us to use the derivation from 2.A.1, equation (2.6).

$$\begin{aligned}E[X - Y | X + Y > c] &= \mu_X - \mu_Y + \lambda \left(\frac{c - \mu_X - \mu_Y}{\sqrt{\sigma_X^2 + \sigma_Y^2 + 2\sigma_{X,Y}}} \right) \left[\frac{\sigma_X^2 - \sigma_Y^2}{\sqrt{\sigma_X^2 + \sigma_Y^2 + 2\sigma_{X,Y}}} \right] \\ &= (1 - s)\mu_r y_1 - s\mu_r y_1 + \lambda \left(\frac{c - (1 - s)\mu_r y_1 - s\mu_r y_1}{\sqrt{((1 - s)y_1 \sigma_r)^2 + (sy_1 \sigma_r)^2}} \right) \left[\frac{((1 - s)y_1 \sigma_r)^2 - (sy_1 \sigma_r)^2}{\sqrt{((1 - s)y_1 \sigma_r)^2 + (sy_1 \sigma_r)^2}} \right] \\ &= \mu_r y_1 (1 - 2s) + \lambda \left(\frac{c - \mu_r y_1}{\sigma_r y_1 \sqrt{(1 - s)^2 + s^2}} \right) \left[\frac{\sigma_r^2 y_1^2 [(1 - s)^2 - s^2]}{\sigma_r y_1 \sqrt{(1 - s)^2 + s^2}} \right] \\ &= \mu_r y_1 (1 - 2s) + \lambda \left(\frac{c - \mu_r y_1}{\sigma_r y_1 \sqrt{(1 - s)^2 + s^2}} \right) \left[\frac{\sigma_r y_1 (1 - 2s)}{\sqrt{(1 - s)^2 + s^2}} \right]\end{aligned}$$

The expression we end up with is given below:

$$E[X - Y | X + Y > c] = (1 - 2s) \left[\mu_r y_1 + \lambda \left(\frac{c - \mu_r y_1}{\sigma_r y_1 \sqrt{(1 - s)^2 + s^2}} \right) \left[\frac{\sigma_r y_1}{\sqrt{(1 - s)^2 + s^2}} \right] \right] \quad (2.3)$$

When $0 < s < 0.5$, the first term, $1 - 2s$, is greater than zero. Inside the brackets, $\mu_r y_1 > 0$ because the mean income in the second period and household income of the first period is assumed to be greater than zero. The Inverse Mills Ratio, $\lambda(\cdot)$ is always greater than zero. And lastly the fraction $\frac{\sigma_r y_1}{\sqrt{(1 - s)^2 + s^2}} > 0$ because $\sigma_r > 0$ and the income is assumed to be greater than zero.

This implies $E[X - Y | X + Y > c] > 0$, proving that the expected return to moving

conditional on moving is larger for men than women for any household with $0 < s < 0.5$.

Proposition 5 *If $\mu_M > \mu_F$ and all households are unitary households, then the expected return to moving (conditional on moving) for men and women is equal for households at $s = 0.5$; i.e., $E[\Delta y_M - \Delta y_F | s = 0.5, \Delta y_M + \Delta y_F > c] = 0$.*

Proof. Note that the expectation, $E[\Delta y_M - \Delta y_F | s = 0.5, \Delta y_M + \Delta y_F > c]$, in this proposition is the same as in 2, but rather than the expression being greater than zero at $0 < s < 0.5$, we want to show that the expression is equal to zero at $s = 0.5$.

Following the same steps to simplify the expectation as in 2, we get equation (2.3) which is reproduced below.

$$E[X - Y | X + Y > c] = (1 - 2s) \left[\mu_r y_1 + \lambda \left(\frac{c - \mu_r y_1}{\sigma_r y_1 \sqrt{(1-s)^2 + s^2}} \right) \left[\frac{\sigma_r y_1}{\sqrt{(1-s)^2 + s^2}} \right] \right]$$

When $s = 0.5$, the first term, $1 - 2s$, is equal to zero which implies $E[X - Y | X + Y > c] = 0$, proving that the expected return to moving conditional on moving is the same for the man and woman for any household with $s = 0.5$.

Proposition 6 *If $\mu_M > \mu_F$ and all households are non-unitary households with $0 < \beta < 1$, then the expected return to moving (conditional on moving), then $E[\Delta y_M - \Delta y_F | s = 0.5, \Delta y_M + \beta \Delta y_F > c] > 0$ with the expectation approaching 0 as β approaches 1 from below.*

Proof. To start, we expand the expectation, $E[\Delta y_M - \Delta y_F | s, \Delta y_M + \beta \Delta y_F > c]$.

$$\begin{aligned} \Delta y_M - \Delta y_F &= (y_{M2} - y_{M1}) - (y_{F2} - y_{F1}) \\ &= \varepsilon_{M2}(1-s)y_1 - \varepsilon_{F2}sy_1 \\ \Delta y_M + \beta \Delta y_F &= (y_{M2} - y_{M1}) + \beta(y_{F2} - y_{F1}) \\ &= (1 + \varepsilon_{M2})(1-s)y_1 - (1-s)y_1 + \beta(1 + \varepsilon_{F2})sy_1 - \beta sy_1 \\ &= \varepsilon_{M2}(1-s)y_1 + \beta \varepsilon_{F2}sy_1 \\ \implies E[\Delta y_M - \Delta y_F | s, \Delta y_M + \beta \Delta y_F > c] &= E[\varepsilon_{M2}(1-s)y_1 - \varepsilon_{F2}sy_1 | s, \varepsilon_{M2}(1-s)y_1 + \beta \varepsilon_{F2}sy_1 > c] \end{aligned}$$

We want to show that when $s = 0.5$, $E[\varepsilon_{M2}(1-s)y_1 - \varepsilon_{F2}sy_1 | s, \varepsilon_{M2}(1-s)y_1 + \beta \varepsilon_{F2}sy_1 > c] > 0$. Using the same substitutions for X and Y as in 2, equation (2.2) at $s = 0.5$, we have $X, Y \sim N(0.5\mu_r y_1, ((0.5y_1\sigma_r)^2))$.

Rewriting the expectation to fit the form, $E[X - Y | X + \beta Y > c]$, and using the results

from 2.A.1, equation (2.7), we plug in our substitutions for X, Y .

$$\begin{aligned}
E[X - Y \mid X + bY > c] &= \mu_X - \mu_Y + \lambda \left(\frac{c - \mu_X - b\mu_Y}{\sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2\sigma_{X,Y}}} \right) \left[\frac{\sigma_X^2 + (b^3 - 2b^2)\sigma_Y^2 + (b-1)\sigma_{X,Y}}{\sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2\sigma_{X,Y}}} \right] \\
&= \lambda \left(\frac{c - 0.5\mu_r y_1 - \beta 0.5\mu_r y_1}{\sqrt{(0.5y_1\sigma_r)^2 + \beta^2(0.5y_1\sigma_r)^2}} \right) \left[\frac{(0.5y_1\sigma_r)^2 + (\beta^3 - 2\beta^2)(0.5y_1\sigma_r)^2}{\sqrt{(0.5y_1\sigma_r)^2 + \beta^2(0.5y_1\sigma_r)^2}} \right] \\
&= \lambda \left(\frac{c - 0.5\mu_r y_1(1 + \beta)}{0.5y_1\sigma_r\sqrt{1 + \beta^2}} \right) \left[\frac{(0.5y_1\sigma_r)^2(1 + \beta^3 - 2\beta^2)}{0.5y_1\sigma_r\sqrt{1 + \beta^2}} \right]
\end{aligned}$$

The expression we end up with at $s = 0.5$ is given below:

$$E[X - Y \mid X + \beta Y > c] = \lambda \left(\frac{c - 0.5\mu_r y_1(1 + \beta)}{0.5y_1\sigma_r\sqrt{1 + \beta^2}} \right) \left[\frac{0.5y_1\sigma_r(1 + \beta^3 - 2\beta^2)}{\sqrt{1 + \beta^2}} \right] \quad (2.4)$$

To prove the proposition, we want to show that the expression above is positive. The Inverse Mills Ratio, $\lambda(\cdot)$, is always greater than zero. And for $0 < \beta < 1$, the numerator in the second term, $0.5y_1\sigma_r(1 + \beta^3 - 2\beta^2)$, is in the open interval $(0, 0.5y_1\sigma_r)$. Because $0.5y_1\sigma_r > 0$, we have shown that $E[X - Y \mid X + \beta Y > c] > 0$, proving that the expected return to moving conditional on moving is the larger for the man and woman for any household with $s = 0.5$ and $0 < \beta < 1$.

Additionally, we want to show that the expectation approaches 0 as β approaches 1. We can do this by taking the limit of the expectation at $s = 0.5$ below:

$$\begin{aligned}
\lim_{\beta \rightarrow 1} E[X - Y \mid X + \beta Y > c] &= \lim_{\beta \rightarrow 1} \lambda \left(\frac{c - 0.5\mu_r y_1(1 + \beta)}{0.5y_1\sigma_r\sqrt{1 + \beta^2}} \right) \left[\frac{0.5y_1\sigma_r(1 + \beta^3 - 2\beta^2)}{\sqrt{1 + \beta^2}} \right] \\
&= \lambda \left(\frac{c - 0.5\mu_r y_1(1 + 1)}{0.5y_1\sigma_r\sqrt{1 + 1^2}} \right) \left[\frac{0.5y_1\sigma_r(1 + 1^3 - 2(1)^2)}{\sqrt{1 + 1^2}} \right] \\
&= \lambda \left(\frac{c - 0.5\mu_r y_1(1 + 1)}{0.5y_1\sigma_r\sqrt{2}} \right) \left[\frac{0.5y_1\sigma_r(0)}{\sqrt{2}} \right] \\
&= 0
\end{aligned}$$

2.A.1 Additional Theoretical Results

In the section below, general derivations are provided based on the following normally distributed random variables.²¹ Let $X \sim N(\mu_X, \sigma_X^2)$, $Y \sim N(\mu_Y, \sigma_Y^2)$, $X + Y \sim N(\mu_X + \mu_Y, \sigma_X^2 + \sigma_Y^2 + 2cov(X, Y))$, and c be a constant.

²¹Some of the results provided in this section are restatements from Heidi Williams' lecture notes on models of self selection available through MIT OpenCourseWare.

$E[X \mid X + Y > c - \mu_X - \mu_Y]$ **with** $(X + Y, X)$ **bivariate normal**

We want to simplify to expression: $E[X \mid X + Y > c - \mu_X - \mu_Y]$. In the first step below, we standardize the expectation (e.g. $\frac{x - \mu_x}{\sigma_x}$ where x is a random variable):

$$\begin{aligned} E[X \mid X + Y > c - \mu_X - \mu_Y] &= \frac{1}{\sigma_X} E \left[\frac{X}{\sigma_X} \mid \frac{X + Y}{\sigma_{X+Y}} > \frac{c - \mu_X - \mu_Y}{\sigma_{X+Y}} \right] \\ E \left[\frac{X}{\sigma_X} \mid \frac{X + Y}{\sigma_{X+Y}} > \frac{c - \mu_X - \mu_Y}{\sigma_{X+Y}} \right] &= E \left[E \left[\frac{X}{\sigma_X} \mid \frac{X + Y}{\sigma_{X+Y}} \right] \mid \frac{X + Y}{\sigma_{X+Y}} > \frac{c - \mu_X - \mu_Y}{\sigma_{X+Y}} \right] \end{aligned}$$

The last line above follows from a version of the law of iterated expectations: for any non-stochastic function $f(\cdot)$ and $X = f(W)$, $E[Y|X] = E[E[Y|X]|X]$.

To simplify the expression further, we want to solve for $E \left[\frac{X}{\sigma_X} \mid \frac{X+Y}{\sigma_{X+Y}} \right]$. Let $s = \frac{X+Y}{\sigma_{X+Y}}$. For simplicity, we assume $\mu_X = \mu_Y = 0$, which would allow and $s \sim N(0, 1)$.

$$\begin{aligned} E \left[\frac{X}{\sigma_X} \mid \frac{X + Y}{\sigma_{X+Y}} \right] &= \frac{1}{\sigma_X} E \left[X \mid \frac{X + Y}{\sigma_{X+Y}} \right] \\ &= \frac{1}{\sigma_X} E[X \mid s] \end{aligned}$$

We need an expression for $E[X \mid s]$, which we can derive using the facts below.

- Given a vector of random variables $X \sim N(\mu, \Sigma)$, then $AX + b \sim N(A\mu + b, A\Sigma A')$. Using this property, because X is normally distributed and $X + Y$ is normally distributed, we know that $\begin{pmatrix} X+Y \\ X \end{pmatrix}$ are jointly normally distributed.
- Given $\begin{pmatrix} X \\ Y \end{pmatrix} \sim N \left(\begin{pmatrix} \mu_X \\ \mu_Y \end{pmatrix}, \begin{pmatrix} \sigma_X^2 & \sigma_{X,Y} \\ \sigma_{X,Y} & \sigma_Y^2 \end{pmatrix} \right)$, then $(Y \mid X = x) \sim N \left(\mu_Y + \rho_{X,Y} \left(\frac{\sigma_Y}{\sigma_X} \right) (x - \mu_X), \sigma_Y^2 (1 - \rho_{X,Y}^2) \right)$. Applying this property to X and $X + Y$, because they are jointly normal, we have $E[X \mid X + Y] = \rho_{X, X+Y} \left(\frac{\sigma_X}{\sigma_{X+Y}} \right) (X + Y) = \frac{\sigma_{X, X+Y}}{\sigma_{X+Y}^2} (X + Y)$.

Adapting those facts to our substitution with s , we have $E[X \mid s] = \rho_{X,s} (\sigma_X / \sigma_s) \cdot s = (\sigma_{X,s} / \sigma_s^2) \cdot s$.

Continuing the substitution,

$$\begin{aligned} E \left[\frac{X}{\sigma_X} \mid \frac{X + Y}{\sigma_{X+Y}} \right] &= \frac{1}{\sigma_X} E[X \mid s] \\ &= \frac{1}{\sigma_X} \frac{\text{cov}(X, s)}{\sigma_s^2} \cdot s \\ &= \frac{1}{\sigma_X} \left[\frac{\text{cov}(X, \frac{X+Y}{\sigma_{X+Y}})}{\sigma_s^2} \right] \cdot s \\ &= \frac{1}{\sigma_X} \frac{\frac{1}{\sigma_{X+Y}} \text{cov}(X, X + Y)}{1} \cdot \frac{X + Y}{\sigma_{X+Y}} \\ &= \frac{\text{cov}(X, X + Y)}{\sigma_X \cdot \sigma_{X+Y}} \frac{X + Y}{\sigma_{X+Y}} \\ &= \rho_{X, X+Y} \frac{X + Y}{\sigma_{X+Y}} \end{aligned}$$

Plugging these results back into the first expression at the beginning of the section:

$$\begin{aligned}
E\left[\frac{X}{\sigma_X} \mid \frac{X+Y}{\sigma_{X+Y}} > \frac{c - \mu_X - \mu_Y}{\sigma_{X+Y}}\right] &= E\left[E\left[\frac{X}{\sigma_X} \mid \frac{X+Y}{\sigma_{X+Y}}\right] \mid \frac{X+Y}{\sigma_{X+Y}} > \frac{c - \mu_X - \mu_Y}{\sigma_{X+Y}}\right] \\
&= E\left[\rho_{X,X+Y} \frac{X+Y}{\sigma_{X+Y}} \mid \frac{X+Y}{\sigma_{X+Y}} > \frac{c - \mu_X - \mu_Y}{\sigma_{X+Y}}\right] \\
&= \rho_{X,X+Y} E\left[\frac{X+Y}{\sigma_{X+Y}} \mid \frac{X+Y}{\sigma_{X+Y}} > \frac{c - \mu_X - \mu_Y}{\sigma_{X+Y}}\right]
\end{aligned}$$

The expectation in the last equation above follows a truncated normal distribution, so we can rewrite it as:

$$E\left[\frac{X}{\sigma_X} \mid \frac{X+Y}{\sigma_{X+Y}} > \frac{c - \mu_X - \mu_Y}{\sigma_{X+Y}}\right] = \rho_{X,X+Y} \frac{\phi\left(\frac{c - \mu_X - \mu_Y}{\sigma_{X+Y}}\right)}{1 - \Phi\left(\frac{c - \mu_X - \mu_Y}{\sigma_{X+Y}}\right)} \quad (2.5)$$

This result will be used to simplify expressions in 2.A.1 and 2.A.1.

$E[X - Y \mid X + Y > c]$ with (X, Y) bivariate normal

We want to calculate $E[X - Y \mid X + Y > c]$ where c is a constant.

$$E[X - Y \mid X + Y > c] = 2E[X \mid X + Y > c] - E[X + Y \mid X + Y > c]$$

We solve for each term separately, starting with the first term: $E[X \mid X + Y > c]$. Redefine $X = \mu_X + \varepsilon_X$ with $\varepsilon_X \sim N(0, \sigma_X^2)$, $Y = \mu_Y + \varepsilon_Y$ with $\varepsilon_Y \sim N(0, \sigma_Y^2)$. It follows that $\varepsilon_X + \varepsilon_Y \sim N(0, \sigma_X^2 + \sigma_Y^2 + 2cov(X, Y))$.

$$\begin{aligned}
E[X \mid X + Y > c] &= E[\mu_X + \varepsilon_X \mid (\mu_X + \varepsilon_X) + (\mu_Y + \varepsilon_Y) > c] \\
&= \mu_X + E[\varepsilon_X \mid \varepsilon_X + \varepsilon_Y > c - \mu_X - \mu_Y] \\
&= \mu_X + \sigma_X E\left[\frac{\varepsilon_X}{\sigma_X} \mid \frac{\varepsilon_X + \varepsilon_Y}{\sqrt{\sigma_X^2 + \sigma_Y^2 + 2cov(X, Y)}} > \frac{c - \mu_X - \mu_Y}{\sqrt{\sigma_X^2 + \sigma_Y^2 + 2cov(X, Y)}}\right]
\end{aligned}$$

To simplify the second term above, we apply the result derived in 2.A.1, equation (2.5).

Let $z = \frac{c - \mu_X - \mu_Y}{\sqrt{\sigma_X^2 + \sigma_Y^2 + 2cov(X, Y)}}$ and $\lambda(z) = \frac{\phi(z)}{1 - \Phi(z)}$.

$$\begin{aligned}
E[X \mid X + Y > c] &= \mu_X + \sigma_X \rho_{\varepsilon_X, \varepsilon_X + \varepsilon_Y} \lambda(z) \\
&= \mu_X + \sigma_X \frac{cov(\varepsilon_X, \varepsilon_X + \varepsilon_Y)}{\sigma_{\varepsilon_X} \cdot \sigma_{\varepsilon_X + \varepsilon_Y}} \lambda(z) \\
&= \mu_X + \sigma_X \frac{var(\varepsilon_X) + cov(\varepsilon_X, \varepsilon_Y)}{\sigma_{\varepsilon_X} \cdot \sigma_{\varepsilon_X + \varepsilon_Y}} \lambda(z) \\
&= \mu_X + \sigma_X \frac{\sigma_X^2 + \sigma_{X,Y}}{\sigma_X \cdot \sqrt{\sigma_X^2 + \sigma_Y^2 + 2cov(X, Y)}} \lambda(z) \\
&= \mu_X + \frac{\sigma_X^2 + \sigma_{X,Y}}{\sqrt{\sigma_X^2 + \sigma_Y^2 + 2cov(X, Y)}} \lambda(z)
\end{aligned}$$

The second term, $E[X + Y|X + Y > c]$, follows a truncated normal distribution which is given by:

$$E[X + Y|X + Y > c] = \mu_X + \mu_Y + \sqrt{\sigma_X^2 + \sigma_Y^2 + 2cov(X, Y)}\lambda(z)$$

Combining the terms together, we get:

$$\begin{aligned} E[X - Y|X + Y > c] &= 2E[X|X + Y > c] - E[X + Y|X + Y > c] \\ &= 2 \left[\mu_X + \frac{\sigma_X^2 + \sigma_{X,Y}}{\sqrt{\sigma_X^2 + \sigma_Y^2 + 2cov(X, Y)}}\lambda(z) \right] - \mu_X - \mu_Y - \sqrt{\sigma_X^2 + \sigma_Y^2 + 2cov(X, Y)}\lambda(z) \\ &= 2\mu_X - \mu_X - \mu_Y + \lambda(z) \left[\frac{2\sigma_X^2 + 2\sigma_{X,Y}}{\sqrt{\sigma_X^2 + \sigma_Y^2 + 2cov(X, Y)}} - \sqrt{\sigma_X^2 + \sigma_Y^2 + 2cov(X, Y)} \right] \\ &= \mu_X - \mu_Y + \lambda(z) \left[\frac{2\sigma_X^2 + 2\sigma_{X,Y} - \sigma_X^2 - \sigma_Y^2 - 2\sigma_{X,Y}}{\sqrt{\sigma_X^2 + \sigma_Y^2 + 2\sigma_{X,Y}}} \right] \end{aligned}$$

The final simplified form for the expression, $E[X - Y|X + Y > c]$, is given below:

$$E[X - Y|X + Y > c] = \mu_X - \mu_Y + \lambda \left(\frac{c - \mu_X - \mu_Y}{\sqrt{\sigma_X^2 + \sigma_Y^2 + 2\sigma_{X,Y}}} \right) \left[\frac{\sigma_X^2 - \sigma_Y^2}{\sqrt{\sigma_X^2 + \sigma_Y^2 + 2\sigma_{X,Y}}} \right] \quad (2.6)$$

$E[X - Y|X + bY > c]$ with (X, Y) bivariate normal

We want to calculate $E[X - Y|X + bY > c]$ where $0 < b < 1$ and c is a constant.

$$E[X - Y|X + bY > c] = 2E[X|X + bY > c] - E[X + bY|X + bY > c] - (1 - b)E[Y|X + bY > c]$$

We solve for each term above separately, starting with the first term: $E[X|X + bY > c]$. Redefine $X = \mu_X + \varepsilon_X$ with $\varepsilon_X \sim N(0, \sigma_X^2)$. Similarly, let $bY = b\mu_Y + \varepsilon_Y$ where $\varepsilon_Y \sim N(0, b^2\sigma_Y^2)$. It follows that $\varepsilon_X + \varepsilon_Y \sim N(0, \sigma_X^2 + b^2\sigma_Y^2 + 2cov(X, Y))$.

$$\begin{aligned} E[X|X + bY > c] &= E[\mu_X + \varepsilon_X | (\mu_X + \varepsilon_X) + (b\mu_Y + \varepsilon_Y) > c] \\ &= \mu_X + E[\varepsilon_X | \varepsilon_X + \varepsilon_Y > c - \mu_X - b\mu_Y] \\ &= \mu_X + \sigma_X E \left[\frac{\varepsilon_X}{\sigma_X} \mid \frac{\varepsilon_X + \varepsilon_Y}{\sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2cov(X, Y)}} > \frac{c - \mu_X - b\mu_Y}{\sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2cov(X, Y)}} \right] \end{aligned}$$

To simplify the second term above, we apply the result derived in 2.A.1, equation (2.5).

As in 2.A.1, we let $z = \frac{c - \mu_X - b\mu_Y}{\sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2cov(X, Y)}}$, $\lambda(z) = \frac{\phi(z)}{1 - \Phi(z)}$, and apply the same steps.

$$\begin{aligned} E[X|X + bY > c] &= \mu_X + \sigma_X \rho_{\varepsilon_X, \varepsilon_X + \varepsilon_Y} \lambda(z) \\ &= \mu_X + \sigma_X \left(\frac{var(\varepsilon_X) + cov(\varepsilon_X, \varepsilon_Y)}{\sigma_{\varepsilon_X} \cdot \sigma_{\varepsilon_X + \varepsilon_Y}} \right) \lambda(z) \\ &= \mu_X + \left(\frac{\sigma_X^2 + \sigma_{X,Y}}{\sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2cov(X, Y)}} \right) \lambda(z) \end{aligned}$$

The second term, $E[X + Y|X + Y > c]$, follows a truncated normal distribution and can be rewritten as:

$$E[X + Y|X + Y > c] = \mu_X + b\mu_Y + \sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2cov(X, Y)} \lambda(z)$$

The third and final term, $E[Y|X + bY > c]$, can be rewritten following a similar derivation to the first term.

$$\begin{aligned} E[Y | X + bY > c] &= \mu_Y + E[\varepsilon_Y | (\mu_X + \varepsilon_X) + (b\mu_Y + \varepsilon_Y) > c] \\ &= \mu_Y + b\sigma_Y E \left[\frac{\varepsilon_Y}{b\sigma_Y} \mid \varepsilon_X + \varepsilon_Y > c - \mu_X - b\mu_Y \right] \\ &= \mu_Y + b\sigma_Y \rho_{\varepsilon_Y, \varepsilon_X + \varepsilon_Y} \lambda(z) \\ &= \mu_Y + b\sigma_Y \frac{cov(\varepsilon_Y, \varepsilon_X + \varepsilon_Y)}{\sigma_{\varepsilon_Y} \cdot \sigma_{\varepsilon_X + \varepsilon_Y}} \lambda(z) \\ &= \mu_Y + b\sigma_Y \frac{var(\varepsilon_Y) + cov(\varepsilon_X, \varepsilon_Y)}{b\sigma_Y \cdot \sigma_{\varepsilon_X + \varepsilon_Y}} \lambda(z) \\ &= \mu_Y + \frac{b^2\sigma_Y^2 + \sigma_{X,Y}}{\sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2\sigma_{X,Y}}} \lambda(z) \end{aligned}$$

Combining all three terms to solve the expression $E[X - Y | X + bY > c]$, we have:

$$\begin{aligned} E[X - Y | X + bY > c] &= \\ &= 2E[X|X + bY > c] - E[X + bY|X + bY > c] - (1 - b)E[Y|X + bY > c] \\ &= 2 \left[\mu_X + \left(\frac{\sigma_X^2 + \sigma_{X,Y}}{\sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2\sigma_{X,Y}}} \right) \lambda(z) \right] - \left[\mu_X + b\mu_Y + \sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2\sigma_{X,Y}} \lambda(z) \right] \\ &\quad - (1 - b) \left[\mu_Y + \frac{b^2\sigma_Y^2 + \sigma_{X,Y}}{\sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2\sigma_{X,Y}}} \lambda(z) \right] \\ &= 2\mu_X - \mu_X - b\mu_Y - \mu_Y + b\mu_Y + \lambda(z) \left[\frac{2\sigma_X^2 + 2\sigma_{X,Y}}{\sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2\sigma_{X,Y}}} - \sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2\sigma_{X,Y}} \right. \\ &\quad \left. - \frac{b^2\sigma_Y^2 + \sigma_{X,Y}}{\sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2\sigma_{X,Y}}} + \frac{b^3\sigma_Y^2 + b\sigma_{X,Y}}{\sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2\sigma_{X,Y}}} \right] \\ &= \mu_X - \mu_Y + \lambda(z) \left[\frac{2\sigma_X^2 + 2\sigma_{X,Y} - \sigma_X^2 - b^2\sigma_Y^2 - 2\sigma_{X,Y} - b^2\sigma_Y^2 - \sigma_{X,Y} + b^3\sigma_Y^2 + b\sigma_{X,Y}}{\sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2\sigma_{X,Y}}} \right] \end{aligned}$$

To summarize, the final derivation is given below:

$$E[X - Y | X + bY > c] = \mu_X - \mu_Y + \lambda \left(\frac{c - \mu_X - b\mu_Y}{\sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2\sigma_{X,Y}}} \right) \left[\frac{\sigma_X^2 + (b^3 - 2b^2)\sigma_Y^2 + (b-1)\sigma_{X,Y}}{\sqrt{\sigma_X^2 + b^2\sigma_Y^2 + 2\sigma_{X,Y}}} \right] \quad (2.7)$$

2.A.2 Model Extensions

Proposition 2 *If $\mu_M > \mu_F$ and all households are unitary households, then the expected return to moving (conditional on moving) for men and women is equal for households at $s = 0.5$; i.e., $E[\Delta y_M - \Delta y_F | s = 0.5, \Delta y_M + \Delta y_F > c] = 0$.*

Proof. Refer to 2.A, Proposition 5.

Corollary 0.1 *Proposition 5 holds in the assortative matching case (i.e., $\rho_{\varepsilon_{M1}, \varepsilon_{F1}} \neq 0$).*

Proof. Recall the substitution for X and Y from equation (2.2) where $X \sim N((1-s)\mu_r y_1, ((1-s)y_1\sigma_r)^2)$ and $Y \sim N(s\mu_r y_1, (s y_1\sigma_r)^2)$. Using this substitution, the expanded form for the expression, $E[\Delta y_M - \Delta y_F | \Delta y_M + \Delta y_F > c]$, is given in Lemma 2, equation (2.3) which is reproduced below.

$$E[X - Y | X + Y > c] = (1 - 2s) \left[\mu_r y_1 + \lambda \left(\frac{c - \mu_r y_1}{\sigma_r y_1 \sqrt{(1-s)^2 + s^2}} \right) \left[\frac{\sigma_r y_1}{\sqrt{(1-s)^2 + s^2}} \right] \right]$$

Notice that X , Y , and $E[X - Y | X + Y > c]$ do not depend on any functional form assumptions on Period 1 income, which is where $\rho_{\varepsilon_{M1}, \varepsilon_{F1}}$ would impact each household member's income. Therefore, assortative matching in the first period will not affect the results and Proposition 5 still holds.

Corollary 0.2 *Proposition 5 holds in the heteroskedasticity case (i.e., $\sigma_M^2 \neq \sigma_F^2$).*

Proof. We can follow the same argument laid out in Proposition 5, Corollary 0.1 looking at the substitutions for X and Y , and referring to the expectation in equation (2.3) above. The variances for X and Y do not depend on Period 1 variance, σ_i^2 for $i = \{M, F\}$, or any functional form assumptions on Period 1 income, so $\sigma_M^2 \neq \sigma_F^2$ would not affect the results and Proposition 5 still holds with heteroskedasticity in the first period.

Proposition 3 *If $\mu_M > \mu_F$ and all households are non-unitary households with $0 < \beta < 1$, then the expected return to moving (conditional on moving), then $E[\Delta y_M - \Delta y_F | s = 0.5, \Delta y_M + \beta \Delta y_F > c] > 0$ with the expectation approaching 0 as β approaches 1 from below.*

Proof. Refer to 2.A, Proposition 6.

Corollary 0.3 *Proposition 6 holds in the assortative matching case (i.e., $\rho_{\varepsilon_M, \varepsilon_F} \neq 0$).*

Proof. From 2.A, Proposition 6, the substitution for X and Y remain identical to equation (2.2). The final expression for $E[\Delta y_M - \Delta y_F | s = 0.5, \Delta y_M + \beta \Delta y_F > c]$ is given in equation (2.4), reproduced below:

$$E[X - Y | X + \beta Y > c] = \lambda \left(\frac{c - 0.5\mu_r y_1(1 + \beta)}{0.5y_1\sigma_r\sqrt{1 + \beta^2}} \right) \left[\frac{0.5y_1\sigma_r(1 + \beta^3 - 2\beta^2)}{\sqrt{1 + \beta^2}} \right]$$

The random variables, X and Y , and the expectation above, do not depend on any functional form of Period 1 income, where $\rho_{\varepsilon_{M1}, \varepsilon_{F1}}$ would impact each household member's income. Therefore, assortative matching in the first period will not affect the results and Proposition 6 still holds.

Corollary 0.4 *Proposition 6 holds in the heteroskedasticity case (i.e., $\sigma_{M1}^2 \neq \sigma_{F1}^2$).*

Proof. As before, we can follow the same argument laid out in Proposition 6, Corollary 0.3 looking at the substitutions for X and Y , and referring to the expectation in equation (2.4) above. Again, the variances for X and Y do not depend on Period 1 variance, σ_i^2 for $i = \{M, F\}$, or any functional form assumptions on Period 1 income, so $\sigma_M^2 \neq \sigma_F^2$ would not affect the results and Proposition 6 still holds with heteroskedasticity in the first period.

2.A.3 Model-Based Simulations

In this section, we numerically simulate the model developed in the main text to estimate how the probability of moving varies with the female share of household income and how the earnings effects of moving vary with the female share of household income. We re-simulate the model under different functional form assumptions and different assumptions on assortative mating. One conclusion from these simulations is that the theoretical results in Foged (2016) are sensitive to functional form assumptions, while the earnings effects (at $s = 0.5$ and for $s < 0.5$ remain robust). This suggests that the potential “U-shaped” pattern of household migration (as a function of the female earnings share) may be a less reliable way to infer the discount households place on income earned by the woman compared to the man.

[Simulation evidence to be added here; available upon request]

2.A.4 Extended Model of Child Penalty

In this section we present an extended version of the model of the child penalty in Andresen and Nix (2022) that incorporates our parameter β that governs the relative weight on income earned by the woman compared to the man. In the baseline Andresen and Nix (2022) model, a couple without children makes a joint hours decision (choosing h_M and h_F) to maximize the following household utility function

$$c + \eta_M \frac{(T - h_M)^{(1-\gamma)}}{1 - \gamma} + \eta_F \frac{(T - h_F)^{(1-\gamma)}}{1 - \gamma}$$

subject to the budget constraint $c \leq w_M h_M + w_F h_F$, where w_M and w_F are the wage rates for the man and woman in the household, T is the total time endowment, eta_M and eta_F are value of leisure parameters that are allowed to vary by gender, and γ determines each individual's labor supply elasticity (which is assumed to be the same for simplicity).

When a couple has a child, the household then makes the following joint hours decision (choosing h_M^C and h_F^C)

$$c + \lambda\theta + \eta_M \frac{(T - h_M)^{(1-\gamma)}}{1 - \gamma} + \eta_F \frac{(T - h_F)^{(1-\gamma)}}{1 - \gamma}$$

subject to the same budget constraint, with $\theta = (1/(1-\kappa)) * (T - h_M^C + T - h_F^C)^{(1-\kappa)}$. Following Andresen et al., the θ parameter is interpreted as the benefit of spending time with children, and λ governs the value to the household of this time investment. (Implicitly, this stylized setup assumes that the household completely substitutes leisure time to child-rearing time after the birth of a child.)

In this setup, the change in income after having a child is defined as the ‘‘child penalty’’ and defined as $(w_i h_i^C - w_i h_i)/(w_i h_i)$ for $i = M, F$. In the simulations reported in the main text, we extend this model in one way which is replacing c in the household utility function with $w_M h_M + \beta * w_F h_F$, and we calibrate the model using the estimated β from the model-based estimation.


We choose $\eta_M = \eta_F = 1$, $\kappa = 0.1$, $\gamma = 0.5$, and we choose the baseline gender wage gap to be $w_F/w_M = 0.895$ in Sweden and $w_F/w_M = 0.82$ in Germany. We then simulate the model for $\lambda = 0$ and $\lambda = 0.25$ at the two different values of β and report the change in earnings for men and women in Table 2.7 in the main text. What this simulation exercise shows is that with no gender differences in preferences for spending time in child-rearing, and a realistic gender earnings gap, the estimated β parameters allow us to account for a majority of the so-called female ‘‘child penalty’’ in both Germany and Sweden. Specifically, the smaller value of β in Germany naturally leads to a larger child penalty because the household is behaving ‘‘as if’’ it places less weight on declines in income by the woman compared to the man following the child's arrival in the household.

2.B Appendix Figures and Tables


2.B.1 Other Employment Measures

Figure 2.B.1: Relationship between Moving and Other Employment Measures


(a) Unemployment Benefits, Germany




(b) Unemployment Benefits, Sweden





(c) Days Employed, Germany



(d) No Employer Connection, Sweden

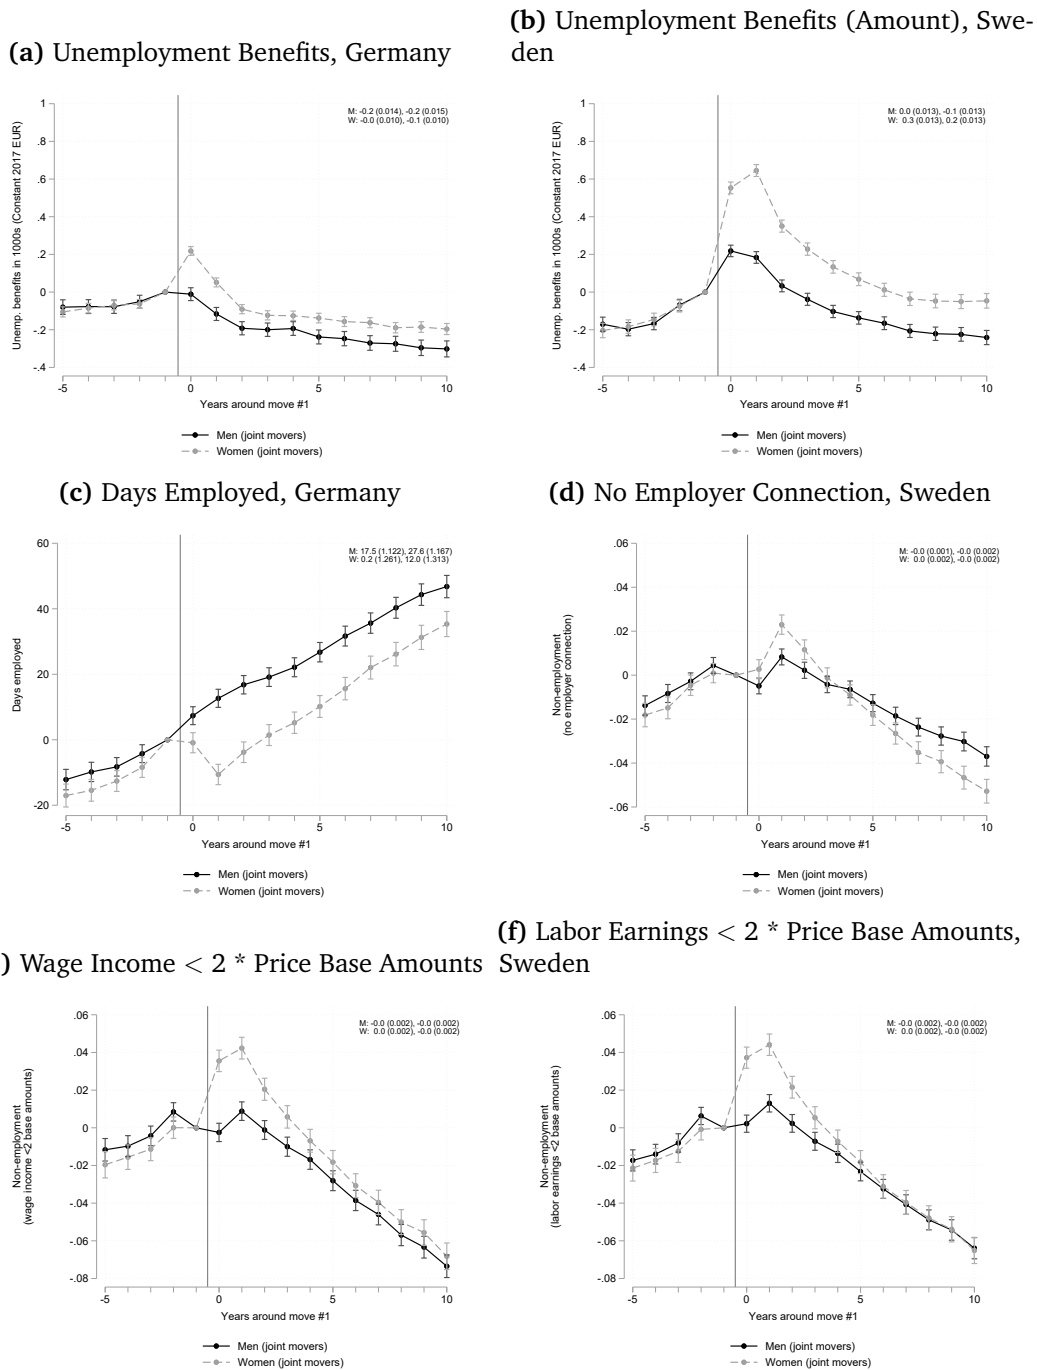


(e) Wage Income < 2 * Price Base Amounts, (f) Labor Earnings < 2 * Price Base Amounts, Sweden



Notes: This figure displays means for different variables in each country from $t - 5$ to $t + 10$ relative to the first move, per gender.

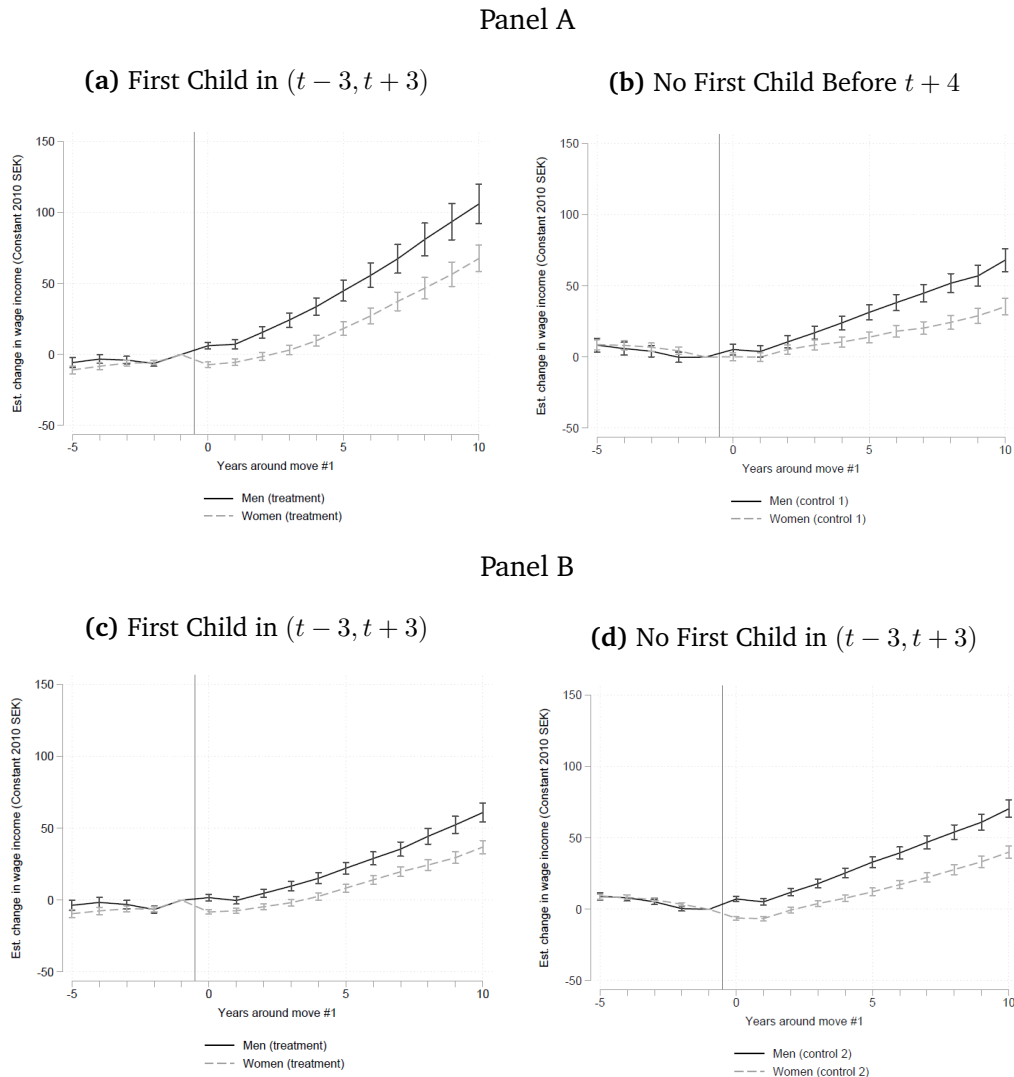
Figure 2.B.2: Event Study Results on Other Measures of Employment



Notes: This figure displays the event study results that estimate the effect of moving on different outcomes in each year relative to the year before the move ($t - 1$). Each point estimate has a corresponding 95% confidence interval calculated using standard errors clustered at the individual level. The regressions are run separately by gender. The coefficients and standard errors (in parentheses) in the upper right corner of each figure are 6 and 11-year averages of the post-move point estimates (from $t = 0$ to $t = 5$ and $t = 10$), in this order, for men (M) and women (W).

2.B.2 Heterogeneity



Figure 2.B.3: Impacts of Move on Wage Income - By Timing of First Joint Child, Sweden





Notes: This figure displays the event study results that estimate the effect of moving on wage income in each year relative to the year before the move ($t - 1$) in Sweden. Each point estimate has a corresponding 95% confidence interval calculated using standard errors clustered at the individual level. The regressions are run separately by gender. The samples of figures (a), (b), and (d) have the following number of observations in $t = 0$: 39,644; 19,720 (33% of the total observations); 56,440 (58% of the total observations). The sample for Figure (c) is the same as for (a). Note that the wage income in this figure is measured with different currency (2010 SEK).

Figure 2.B.4: Impact of Move on Wage Income – By Gender-specific Predicted Female Share of HH Income

(a) Female Share of HH Income < 50%, Ger- (b) Female Share of HH Income < 50%, Swe-
 many den



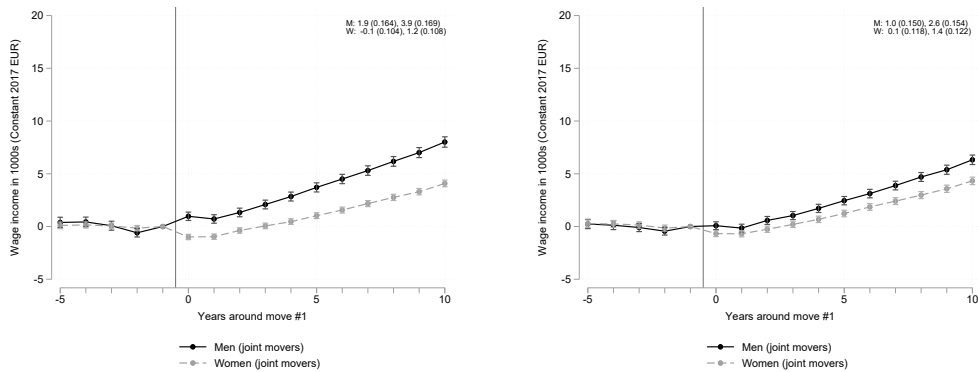
(c) Female Share of HH Income ≥ 50%, Ger- (d) Female Share of HH Income ≥ 50%, Swe-
 many den



Notes: This figure displays the event study results that estimate the effect of moving on wage income in each year relative to the year before the move ($t - 1$). Each point estimate has a corresponding 95% confidence interval calculated using standard errors clustered at the individual level. The regressions are run separately by gender. The coefficients and standard errors (in parentheses) in the upper right corner of each figure are 6 and 11-year averages of the post-move point estimates (from $t = 0$ to $t = 5$ and $t = 10$, in this order, for men (M) and women (W)). Predicted earnings share are calculated regressing log individual income on experience indicators, education level interacted with field of study, and an indicator on having a child under 19 years old.

Figure 2.B.5: Impact of Move on Wage Income – By Predicted Female Share of HH Income, Median

(a) Female Share of HH Income < Median, Sweden (b) Female Share of HH Income \geq Median, Sweden



Notes: This figure displays the event study results that estimate the effect of moving on wage income in each year relative to the year before the move ($t - 1$). Each point estimate has a corresponding 95% confidence interval calculated using standard errors clustered at the individual level. The regressions are run separately by gender. The coefficients and standard errors (in parentheses) in the upper right corner of each figure are 6 and 11-year averages of the post-move point estimates (from $t = 0$ to $t = 5$ and $t = 10$), in this order, for men (M) and women (W). Predicted earnings share are calculated regressing log individual income on experience indicators, education level interacted with field of study, and an indicator on having a child under 19 years old. We do not have these results for Germany yet. The median female share of HH income is 48%.

2.B.3 Predicted Income Methodology and Results

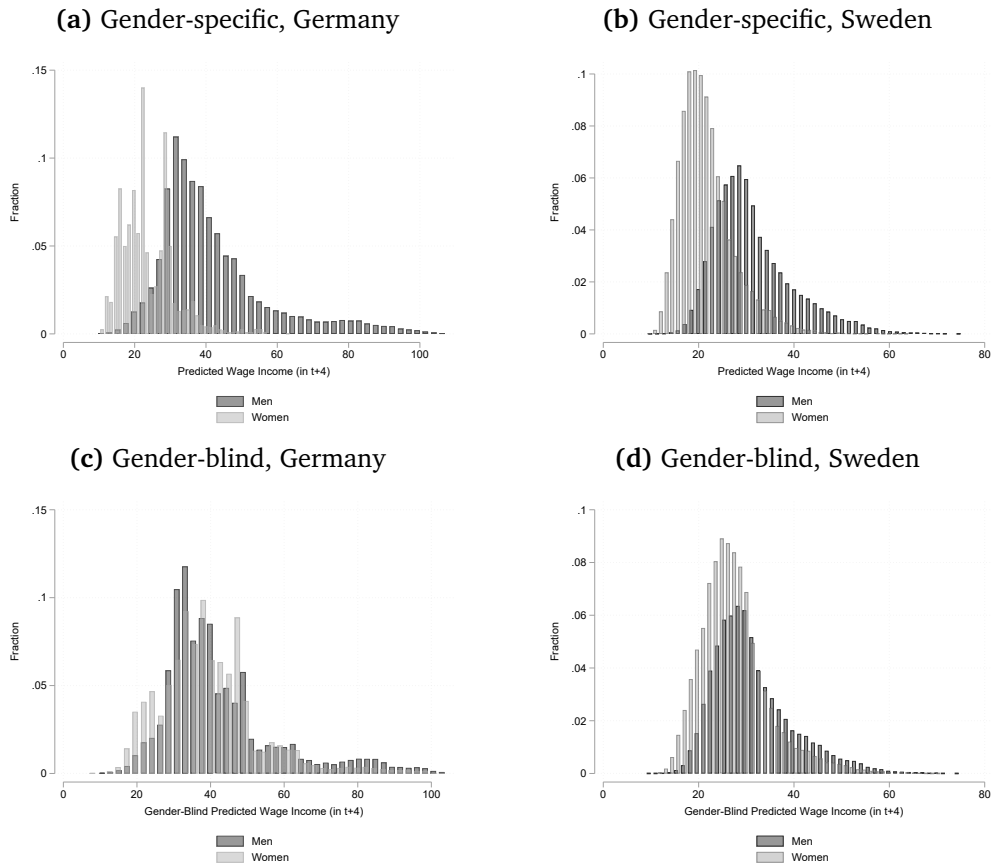
We use the following earnings prediction model:

```
reghdfe lnwageinc i.expproxy, absorb(i.child18 i.lvlfield3 i.year,  
savefe), resid
```

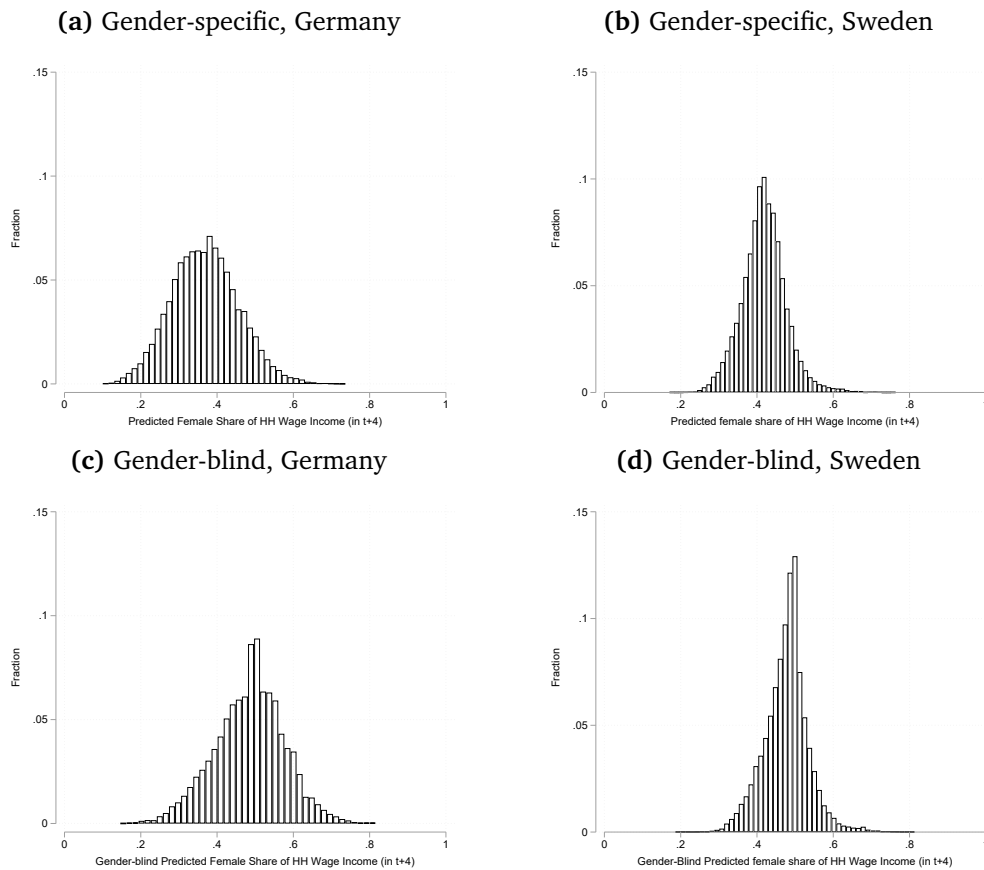
which controls for potential experience, number of children, college major (interacted with highest level of education), and year. In Germany, we do not have college major information so we replace with the highest level of education (three education categories: high school or less, vocational training, some college or more).

We estimate the model in both countries using a 1990-2017 panel with a sample of the population aged 25–54, dropping the individuals with a wage income below 2 price base amounts (which is our preferred proxy for non-employment), and we experimented with alternative models that included additional interactions between level of education.

In the baseline analysis, we focus on *gender-blind* predictions so that the regression model above is run on men and women together. We also report results using *gender-specific* predictions where the regression model above is run on men and women separately.

Figure 2.B.6: Predicted Wage Income, Movers

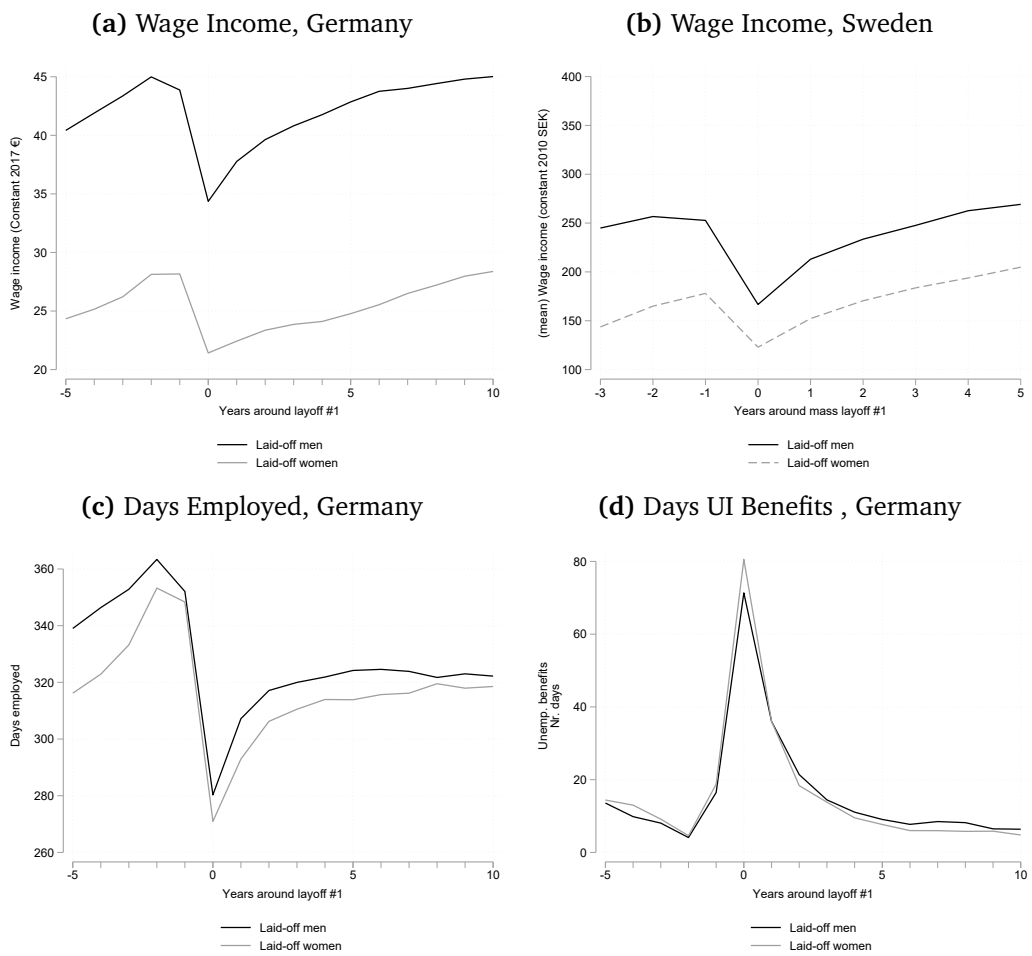
Notes: This figure displays histograms of predicted wage income by gender for each country on the movers sample. Predicted earnings share are calculated regressing log individual income on experience indicators, education level interacted with field of study, and an indicator on having a child under 19 years old. Gender-blind predicted earnings are calculated by regressing men's log individual income on experience indicators and education level interacted with field of study, in a way that men and women with the same covariates have the same predicted wage income.

Figure 2.B.7: Predicted Female Share of HH Income, Movers

Notes: This figure displays histograms of predicted female share of household income by country on the movers sample. Predicted earnings share are calculated regressing log individual income on experience indicators, education level interacted with field of study, and an indicator on having a child under 19 years old. Gender-blind predicted earnings are calculated regressing men's log individual income on experience indicators and education level interacted with field of study, in a way that men and women with the same covariates have the same predicted wage income.

2.B.4 Descriptive Figures for Layoffs

Figure 2.B.8: Relationship between Layoffs and Labor Earnings and Employment



Notes: This figure displays means for different variables in each country from $t - 5$ to $t + 10$ relative to the first layoff event, per gender.

Table 2.B.1: Stepwise Restrictions to Layoffs Sample, Sweden

Sample Nr	Sample restriction	# workplace IDs	# employees
0	Workplace restrictions only (including non-laid off individuals)	1,147	151,150
1	Excl. people staying at $t - 1$ workplace	1,147	104,486
2	Balanced sample (in LISA $t - 3$ to $t + 5$)	1,147	98,473
3	Age restriction 18-65	1,147	97,022
4	Only including 1st layoff	1,147	97,022
5	Excl. ind. working at $t - 1$ workplace in $t = 1/5$	1,147	93,436
6	Requiring tenure ($t - 2/t - 1$)	1,147	68,693
7	Requiring labor market attachment (wage income > 2 pba in $t - 1$)	1,145	57,945
8a	Keeping only married/cohabiting couples	1,117	32,159
8b	8a + w/ unemployment benefits in t	721	3,465
9a	Keeping only marrieds/cohabs, ages 25-45	1,087	16,164
9b	9a + w/ unemployment benefits in t	584	1,996

Notes: This table displays the number of observations for each step in the restrictions applied to the layoffs sample in Sweden.

Alive and Kicking? Short-Term Health Effects of a Physician Strike in Germany

joint work with Daniel Avdic, Martin Karlsson and Nina Schwarz

Chapter Abstract

We study the effects of a physician strike in German hospitals in 2006 on patient mortality. Leveraging a comprehensive dataset encompassing all hospital admissions in Germany and employing digitised records of strike participation, we estimate a difference-in-differences model to discern the causal effects of the strike. Our estimation results reveal a substantial decrease in hospital admissions during the strike period, whereas effects on hospital mortality are mostly driven by patient selection. To support this claim, we further show that emergency cases and more fragile patients, who were unable to substitute their immediate care needs, were more likely to be present in hospital during this period. Hence, in contrast to most other related studies, our results suggest that short term interruptions in access to healthcare may not have dramatic effects on healthcare quality provided that rationing of care by patient severity is carried out.

3.1 Introduction

Labour disruptions by physicians are often seen as controversial because of physicians' obligation towards their patients (Metcalf et al., 2015). The International Labour Organisation (ILO) defines hospitals as an 'essential service' where a strike could in principle be prohibited due to a threat to life, personal safety or health (cf. International Labour Office, 2006). Even temporary disruptions to access to healthcare providers may have important effects on healthcare quality, such as emergency care. However, despite these concerns, physician strikes hit many countries, such as France in 2019, England in 2015, and Portugal in 2012 and 2019 and Germany in 2006 and 2020.

This paper analyses the short-term effects of a nationwide physician strike in Germany on inpatient mortality. The event took place in 2006 when physicians in tertiary hospitals took industrial action for three months between March and June.¹ It was one of the longest labour disruptions in German history and at the peak about 2/3 all physicians employed in tertiary hospitals participated. To empirically study the effects of the strike, we use rich administrative data on all hospital admissions and deaths in Germany for the years 2000-2008 to compare outcomes in strike- and non-strike impacted hospitals over time in a difference-in-differences model. The unprecedented and unanticipated nature of the strike provides a context where causal effects of the disruption on healthcare supply can be plausibly identified and quantified.

The extent of the adverse effects of a physician strike on patient health is not given a priori. On the one hand, labour disruptions may reduce quality of healthcare due to decreased geographical access, higher workload among healthcare staff and crowding effects. Avdic (2016) shows that an increased distance to hospitals due to closures of emergency departments resulted in a lower probability of surviving an acute myocardial infarction in the short run, due to an increase in the risk of out of hospital mortality. Piérard (2014), Lin (2014) and Aiken et al. (2002) show that the staff-to-patient ratio is a determinant of patients' health outcomes. On the other hand, the adverse

¹The majority of German tertiary hospitals are organised on federal state level and thus provide medical care to all publicly as well as privately insured individuals. They incorporate medical schools and are also involved in research and teaching activities.

effects of a strike on healthcare quality can be mitigated through rationing care and triaging patients based on urgency. For example, hospitals can cancel or postpone elective surgeries, which could even lead to lower in-hospital mortality rates and concomitant complications in the short run. Cunningham et al. (2008) and Metcalfe et al. (2015) systematically review the medical literature on physician strikes and conclude that at least the withdrawal of medical services for a short period of time does not seem to increase mortality for the majority of labour disputes. A potential explanation is that in most cases some kind of emergency care is still available. However, most of the reviewed articles only present suggestive evidence and no causal analysis.

Our difference-in-difference estimates suggest that the strike led to a sharp reduction (12%) in hospital admissions in striking hospitals during the months the strike was ongoing compared to non-striking hospitals. Emergency admission rates increased by 5% and length of stay increased by 0.2 days. Estimates not adjusted for case-mix also indicate an increase in in-hospital mortality by 9-10%. Due to the potential endogeneity of patient composition because of care rationing, it is crucial to control for patient case-mix when studying strike effects in healthcare. To do this in the presence of a very large set of covariates representing case-mix, we implement a double selection lasso model (Belloni et al., 2014). Once we adjust for exogenous patient characteristics, the coefficient on mortality is reduced by half. Furthermore, our results indicate some spillover effects to nearby hospitals, whereas there are no signs of a post-strike catch-up effects on admissions. We conclude that healthier patients at striking hospitals likely avoided care or were triaged, so that the patient composition during the strike had a higher share of 'frail' patients with higher underlying mortality risk.

We contribute to the scant literature on the causal effects of strikes in healthcare facilities. To our knowledge, only two papers have investigated the impact of physician strikes. Costa (2019) investigates labour disruptions caused by physicians in Portugal. He investigates the effects of nurses', physicians' and diagnostic and therapeutic technicians' (DTT) strikes on health outcomes. While no effects for nurses' and DTT strikes can be found, physician strikes seem to have increased in-hospital mortality by about 8%. Furthermore, he finds a shift from inpatient to outpatient services and an increase in complications around birth. Stoye and Warner (2023) analyse the

impact of a number of short strikes of junior doctors in the United Kingdom in 2016. Even though our research question is related, our design is very different from those two papers. Whereas they consider short and repeated disruptions caused by strikes, we study the impact of a prolonged strike that affected entire hospitals over a period of three months. Therefore, the shock we consider represents a more serious challenge to the health care system, but it also entails more time for the system to respond to the challenge. In this sense, our analysis provides a very strong test of the resilience of the system in the presence of a major unexpected challenge.

In terms of strikes of other healthcare workers, Gruber and Kleiner (2012) analyse the health effects of nurses' strikes in New York between 1984-2004 and observe that strikes in hospitals lead to lower quality of medical care. They find an 18.3% increase in in-hospital mortality and an increase of 5.7% in 30-day readmission for patients admitted during a strike. Furthermore, Kronborg et al. (2016) investigate a nurse strike in Denmark in 2008. They look at maternal and infant health and conclude that there was a decrease in prenatal midwife consultations, in length of stay after birth and in the number of home visits caused by the strike. In a follow-up study (Hirani et al., 2019) the same strike is used to investigate the impact of the timing of nurse home visits for newborns. Having missed earlier nurse visits due to the strike led to more regular and urgent general practitioner contacts in the first years of life compared to having missed later nurse visits. Furthermore, Friedman and Keats (2019) investigate effects on infant and neonatal health. Their conclusion is that babies born during a health worker strike in Kenya are more likely to die if they were born in a month with strike occurrence. An additional 15 deaths per 1,000 live births is observed one week and one month after birth, which is an increase of 54-68%. We show that the impact of strikes among physicians may generate different impacts on care quality compared to nurses and midwife strikes.

The remainder of the paper proceeds as follows. The next section describes the setting. Section 3 gives information on the different datasets used in the analysis. Section 4 describes our empirical strategy and section 5 presents the results on the effects of the strike. The last section concludes.

3.2 Background

3.2.1 The German Hospital Sector

The German hospital sector treats about 19 million patients every year and involves approximately 2,000 hospitals with 500,000 beds (Federal Statistical Office, 2018). Hospitals can be owned by public, private or non-profit organisations. Public hospitals can be separated further into hospitals run by the municipalities, the 16 federal states, the state and hospitals related to the statutory accident insurance (Goepfert and Conrad, 2013). The majority of the 36 university hospitals are run by the federal states. They are linked to the medical faculties of the universities and are thus also involved in research and teaching activities.

Roughly 10% of the 19 million patients treated in hospitals are admitted to university hospitals (Federal Statistical Office, 2018). Although they make up less than two percent of the hospitals, they treat around 10% of patients and hold about 20% of the physicians working in hospitals.² As there is free choice of healthcare providers, patients can freely choose the hospital they want to be treated in and physicians in the primary or ambulatory care sector are not obliged to transfer patients to a specific hospital (de Cruppé and Geraedts, 2017). Treatment costs are usually covered by a patient's health insurance, as health insurance is mandatory in Germany and all legal residents have universal coverage by either public or private health insurance (Blümel and Busse, 2015).³

3.2.2 Marburger Bund

The *Marburger Bund* is a professional organisation and the trade union for employed physicians in Germany, covering about 70% of all physicians employed in hospitals (Greef, 2012). Prior to the initiation of the physician strike in 2006, it has mainly acted as a professional organisation and was repre-

²Table 3.A.1 in 3.A shows how the number of physicians and other health care workers in hospitals and university hospitals evolved between 2000 and 2010.

³The majority of patients (88%) are publicly insured and only about 11% are covered by a private health insurance (Bundesministerium für Gesundheit, 2019). The data used in this analysis covers privately and publicly insured patients. However, it is not possible to distinguish between privately and publicly insured.

sented during collective bargaining by the *Vereinte Dienstleistungsgewerkschaft (ver.di)*, a multi-branch trade union (Greef, 2012).

Along with the negotiations on a new collective bargaining agreement for public services in 2005, the *Marburger Bund* demanded better working conditions for physicians working in hospitals (Marburger Bund Bundesverband, 2012). As *ver.di* was not willing to negotiate on an own collective bargaining agreement for physicians, the *Marburger Bund* decided to terminate the collective partnership with *ver.di* and become the first and only trade union for physicians.

The *Marburger Bund* and the employer's association officially started collective bargaining for hospitals on federal state level⁴ on October 12, 2005 (Greef, 2012). As they were not able to reach a consent, the biggest-ever physician strike in German history started on March 16, 2006 and went on until June 16, 2006, when finally a collective bargaining agreement on federal state level was achieved. Table 3.1 presents a chronology of events starting in 2005. The table also presents information on the strike in hospitals on the municipality level. However, the focus in this paper will be on the strike at the federal state level, as university hospitals are organised on the federal state level.

Table 3.1: Chronology of Events

Date	Event
14 Apr 2005	<i>ver.di</i> and employer's association start bargaining (federal)
09 Sep 2005	MB terminates collective partnership with <i>ver.di</i>
12 Oct 2005	MB and employer's association start bargaining (federal)
09 Mar 2006	MB and employer's association start bargaining (local)
16 Mar 2006	Strike begins (federal)
16 Jun 2006	MB and employer's association reach agreement for physicians (federal)
26 Jun 2006	Strike begins (local)
17 Aug 2006	MB and employer's association reach agreement for physicians (local)

Notes: MB = Marburger Bund; *ver.di* = Vereinte Dienstleistungsgewerkschaft; Sources: Greef (2012) and Martens (2008).

⁴The *Marburger Bund* had to negotiate with two different employer's associations – the *Tarifgemeinschaft deutscher Länder (TdL)* which is responsible for public hospitals on the federal state level and the *Vereinigung kommunaler Arbeitgeberverbände (VKA)* which is responsible for public hospitals on the municipality level. The focus in this paper will be on the strike at the federal state level, as university hospitals are organised on federal state level.

3.2.3 Physician Strike

The strike studied in this paper took place in the majority of German university and federal hospitals between March and June 2006⁵. It was one of the longest labour disruptions in German history (Martens, 2008). At first, these labour disruptions were restricted to a few university and federal hospitals and to one or two days per week, but soon the strikes expanded to other university hospitals and lasted longer and longer; sometimes even a whole week. At the peak of the strike, about 2/3 of all physicians in these hospitals participated.


In general, emergency, intensive care, and maternity units were generally excluded from the strike. This was ascertained by emergency agreements, signed by the university hospitals and the *Marburger Bund*. However, elective treatments and other non-urgent cases were affected by the strike.

3.3 Data

3.3.1 Strike Data

One challenge with studying the effects of the 2006 strike is that there exist no reliable administrative data which could be used for the analysis (Dribbusch, 2018). Therefore, we collected information on the strike from newspaper articles published in 2006 to obtain information on the exact date of the start and end of the labour disruptions for each university hospital. In total, physicians from 24 out of 33 university hospitals went on strike; nine of these participated from the first day onwards, while the other participating hospitals joined a few weeks later. The remaining nine university hospitals were not part of the collective bargaining agreement with the TdL in early 2006 and therefore did not participate in the labour disruptions. Figure 3.1 shows the location of treatment and control university hospitals in Germany.

⁵There were 33 university hospitals in total, of which 24 participated in the strike. The university hospitals in Berlin, Frankfurt, Gießen and Marburg did not participate, as Berlin and Hesse are not part of the employer's association. Berlin has been excluded from the employer's association in 1994 and Hesse left in 2004 (Greef, 2012). Hamburg and Schleswig-Holstein (Kiel and Lübeck) did not take part as they had reached a collective agreement before the strike took place. The University hospitals in Homburg, Bochum, Mannheim and Wuppertal did not participate in the physician strike for other reasons.

Figure 3.1: Location of Treatment and Control Hospitals

Notes: This figure shows the location of treatment and control hospitals.

3.3.2 Hospital Level Data

Hospital level data on admissions for the years 2000 to 2008 was obtained from the German Federal Statistical Office (FDZ der Statistischen Ämter des Bundes und der Länder, 2008a). It contains individual level information on each hospital admission, including date of admission, length of stay, mortality and surgery indicators, 3-digit ICD-diagnosis and patient characteristics such as age, gender, and place of residence. Hospital level characteristics such as number of physicians and nurses are available on a yearly basis and used for balancing tests (see Section 3.4). Since German hospitals are required to provide this information by law, the data is complete and of high quality (see KHStatV, 1990).

The admission data is aggregated on day of admission and hospital level. For the main regressions on the hospital level, we keep only university hospitals as they provide a balanced sample whereas other hospitals could have been exposed to the strike on the municipality level starting in June 2006,

for which we do not have strike data. The final sample consists of 143,744 observations. Table 3.2 presents descriptive statistics for the main outcome variables, treatment indicator and patient characteristics by treatment status of the hospital over all observation years.

Table 3.2: Descriptive Statistics - Hospital Data

	Treatment Hospital			Control Hospital		
Female	0.49	0.05	84,986	0.49	0.07	58,758
Age	47.38	3.89	84,986	48.99	5.28	58,758
Cases per day per hospital	159.28	69.88	84,986	156.27	111.23	58,758
Log Cases per day per hospital	4.95	0.54	84,986	4.79	0.81	58,758
Length of Stay	8.12	1.59	84,986	8.26	2.13	58,758
Surgeries per 100 admissions	37.06	22.17	84,986	33.57	22.03	58,758
Log Deaths per day per hospital	0.83	0.59	68,985	0.92	0.66	38,410
Death rate per 100 admissions	1.67	1.45	84,986	1.94	2.17	58,758
Death rate within 10 days after admission per 100 admissions	0.91	1.11	84,986	1.08	1.64	58,758
Emergencies per 100 admissions	27.65	12.18	84,648	30.03	13.22	58,508
Urgency of Treatment 0-1	0.33	0.09	84,648	0.35	0.09	58,508

Notes: Descriptive statistics by treatment status of the hospital (weighted by pre-strike cases). Variable definitions are available in appendix section 3.B. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

The primary outcome variable in this analysis is the mortality rate (patient deaths per 100 admissions) in each hospital on each day. As suggested by Gruber and Kleiner (2012), we also investigate deaths occurring within 10 days after the day of admission as this outcome limits bias from strike-induced changes in length of stay. Furthermore, we estimate strike effects on length of stay, surgery rate (surgeries per 100 admissions), an ambulatory care sensitive (ACS) indicator based on Sundmacher et al. (2015)⁶ and an emergency and urgency indicator based on the classification by Krämer et al. (2019)⁷. The latter outcome variables are informative of the extent to which the strike led to any changes of the patient composition in hospitals.

Exposure to the strike is defined as being admitted to a striking hospital during the strike period. The variable takes on the values zero or one. This

⁶The ACS indicator helps to identify hospitalisations which could have been avoided by effective and timely ambulatory care (Sundmacher et al., 2015). The analysis in this paper investigates whether the strike changed the ACS rate of a hospital by transferring patients to ambulatory care units outside of the hospital.

⁷Krämer et al. (2019) developed a diagnosis-based classification to determine whether a hospitalisation is elective or an emergency. They used machine learning and several predictor variables to classify the different diagnosis. A hospitalisation is defined as an emergency if their urgency indicator is above 0.5.

variable does not capture patients that already were hospitalised when the strike began. We therefore use an alternative treatment indicator in a robustness check, defined as the share of patients on a specific day that were affected by the strike at any time during their hospital stay (see Section 3.5.3).

3.3.3 Mortality Census

Mortality data was also obtained from the German Federal Statistical Office (FDZ der Statistischen Ämter des Bundes und der Länder, 2008b) and includes all deaths (both in and out of hospitals) and causes of death (3 digit ICD-Codes) between 2000 and 2008. We use district⁸ of residence and date of death to link and merge the mortality data with the information on strikes we collected and then we collapse the data to the date of death and district level. Treatment assignment is based on district of residence, as the data does not allow for computing the exact distance between place of residence and striking hospitals.

Our final analysis data sample consists of approximately 1.3 million district-date cells which we use as observations in our regressions. The district mortality data is useful to study effects on mortality outside of hospitals. For example, it could be that patients affected by the strike died outside of hospitals if they were discharged too early or failed to be admitted in time for a critical condition. As our data covers the post-strike period, we are also able to study any lagged effects of the strike to study, for example, net mortality effects of the strike. Descriptive statistics for the mortality data are presented in Table 3.3.

⁸The district (*Kreis*) is the intermediate administrative level between municipality (*Gemeinde*) and federal state (*Bundesland*). These units have responsibilities in the domains of public transport, road construction, and the construction of hospitals. There were in total 323 of them in 2006.

Table 3.3: Descriptive Statistics - Mortality Data

	Treatment Community N = 77,477		Control Community N = 32,880	
	Mean	SD	Mean	SD
Treatment Variable	0.03	0.16	0.00	0.00
Female	0.55	0.23	0.54	0.18
Married	0.38	0.23	0.38	0.18
Age at death	76.12	7.19	75.86	5.41
Log Deaths	1.92	0.90	2.46	1.01

Notes: Descriptive statistics by treatment status of the community (weighted by pre-strike cases). Variable definitions are available in appendix section 3.B. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008b), own calculations.

3.4 Econometric Strategy

We apply a difference-in-differences (DID) empirical design to estimate the causal effects of the 2006 strike. To this end, we use the nine non-striking university hospitals as a control group for the 24 striking university hospitals to estimate group-specific changes in mortality and other outcomes of interest before, during and after the strike occurred. Whereas the two groups may exhibit systematic differences in size and staffing levels, they are likely to be exposed to the same general trends in terms of admissions and case-mix. Therefore, assuming that in the absence of the strike, the two groups would follow a parallel trend seems reasonable, and in this case the DID specification would identify the average treatment effect on the treated. Below, we carefully assess the plausibility of this assumption and choose our empirical specification based on the balance of some key hospital characteristics. We also consider an alternative control group in robustness checks. Subsequent analyses, described in detail below, will explore underlying channels for our main results, including changes in patient composition and spillover effects in non-striking hospitals.

3.4.1 Empirical Model

Our main regression specification is defined by the following equation:

$$y_{it} = \beta_0 + \beta_1 S_{it} + \mu_i + \nu_t + H_i \cdot t + \epsilon_{it} \quad (3.1)$$

where y_{it} is one of the outcome variables listed in Table 3.2 for hospital i on day t , S_{it} is a strike indicator which is 1 for strike-hit hospitals during the strike period and 0 otherwise, μ_i are hospital fixed effects, ν_t are date of admission fixed effects and $H_i \cdot t$ are hospital-specific linear time trends. Date of admission fixed effects include fixed effects for year, week, and day of week. The analysis exploits regional and temporal variation in strike exposure and hospital fixed effects can control for time-invariant hospital heterogeneity, such as average case volume, number of beds and share of specialists. Under the assumption that treatment and control hospitals would have followed a common trend in the absence of the strike, conditional on the included control variables, the parameter β_1 measures the average treatment effect on the treated (ATT) of the strike on the outcome of interest. All regressions are weighted by the baseline average admissions per day in each hospital and standard errors are clustered on hospital level.⁹

In addition to in-hospital mortality, we also investigate aggregate district mortality effects by estimating the following regression model:

$$y_{jt} = \beta_0 + \beta_1 S_{jt} + \mu_j + \nu_t + C_j \cdot t + \gamma X_{jt} + \epsilon_{jt} \quad (3.2)$$

where y_{jt} is the logarithmised number of deaths occurring on day t in district j , S_{jt} is a strike indicator variable equal to one for districts with a strike-hit hospital during the strike period and zero otherwise, μ_j are district fixed effects, ν_t are date fixed effects and $C_j \cdot t$ are district specific time trends. Date fixed effects include year fixed effects, week fixed effects and day of week fixed effects. Finally, X_{jt} is a vector of time-varying district-level patient

⁹Although there are only 33 university hospitals in total, some of these are based in multiple locations. We treat each location as a separate entity meaning that there are 44 clusters in total. As number of admissions might be affected by the treatment, we also present unweighted estimates for the main outcome variables in Table 3.A.2 in 3.A. Point estimates are similar to the weighted estimates. Furthermore, we implement case-mix adjustment in Section 3.4.2

characteristics, including the share of females, and the share of married individuals. The coefficient of interest is β_1 , measuring the effect of the strike on the log number of deaths in districts with a striking hospital.

Table 3.4 presents hospital balancing tests to study the comparability between treated and untreated hospitals in our sample.¹⁰

Table 3.4: Covariate Balancing Tests

VARIABLE	N	Mean	(1)	(2)	(3)
N Beds	395	1,130.48	432.597** (191.028)	-22.986 (27.176)	-5.064 (15.843)
Share ICU Beds	395	0.08	0.020** (0.009)	0.001 (0.002)	0.002 (0.003)
N Physicians	395	573.50	337.988*** (106.905)	37.294 (23.928)	8.034 (20.568)
Share Female Physicians	395	0.33	-0.003 (0.017)	0.001 (0.008)	-0.004 (0.011)
N Physicians in Charge	395	40.74	15.316 (12.246)	-2.302 (4.399)	-3.242 (3.192)
N Assistant Physicians	395	433.12	227.567** (87.075)	20.984 (17.948)	10.834 (11.049)
N Non-Physician Employees	395	3,294.67	1972.984*** (560.346)	202.625* (108.468)	44.069 (104.998)
Share Female Non-Physician Employees	395	0.77	-0.010 (0.013)	-0.008* (0.004)	-0.002 (0.003)
N Trainee Positions	395	352.22	134.706** (66.874)	-20.573 (19.103)	-22.718 (17.253)
N Childbirth	395	1,167.57	508.146 (413.760)	51.139 (50.383)	9.083 (38.120)
Share C-sections	306	0.36	0.000 (0.037)	-0.021 (0.025)	-0.008 (0.019)
Year FE			✓	✓	✓
Hospital FE				✓	✓
Hospital Specific Trends					✓

Notes: Robust standard errors clustered at the hospital level are reported in parenthesis. Significance levels: * 0.10 ** 0.05 *** 0.01. Regressions successively include year fixed effects, hospital fixed effects and hospital specific time trends. In comparison to the main specification the treatment period is defined as 2006 as data was only available on a yearly level. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

To study this, we regress each of the variables reported in the table on a set of control variables and the treatment variable indicator. The point estimates from column (1), where we only include year fixed effects as controls, suggest that treated hospitals are on average larger with greater number of beds, more staff and a higher share of intensive care unit beds.

¹⁰As the data is only available yearly, we coded the entire year 2006 as treatment period.

However, once we additionally control for hospital fixed effects and hospital specific trends, the test suggests that the two groups are well-balanced on observable characteristics.

3.4.2 Patient Selection

There are at least two challenges to the identification of causal effects in our setting, which might affect the validity or interpretation of any results derived from our regression model. First, it could be that there are underlying trends in the outcome variables that might have caused the strike. Due to the unprecedented nature of the strike, we argue that it is unlikely that patients were anticipating such radical labour disruptions. Nevertheless, we address this concern by presenting event study estimates for the main outcome variables and test for linear pre-trends.

A second issue regarding the estimation of strike effects on mortality is that the strike might have changed the patient composition. While emergency care was still available, elective surgeries and other non-emergent cases might have been affected. If emergency care patients have worse health outcomes than non-emergency care patients, this change in the case-mix will contribute to our DID estimates. Whereas the estimated effect captures the causal effect of the strike on the hospital-level mortality rates, it is not the causal effect that is relevant from a policy perspective. In order to infer whether the quality of care was compromised during the strike, we need to base estimates of mortality effects on a comparable population of patients. We address this concern in two ways: first, we seek to control for patient composition in the empirical analysis, and second, we estimate effects of the strike on district mortality rates.

Combining the various patient and case characteristics that are available in the data – age, gender, ICD codes, indicators of urgency, emergency and deferability – there are 337 specific covariates that may be considered. In order to avoid overfitting, we rely on a double selection Lasso model, where covariates that predict either our main treatment indicator, or the mortality outcome, are included in the specification (Belloni et al., 2014). For example, our lasso selection specification eventually selected 163 covariates for the analysis of the mortality rate. We add them as control variables to the main

regression specification.

3.4.3 Analysis of Spillover Effects

If admissions in striking hospitals decrease and only more fragile patients are admitted, healthier patients may decide to either delay hospitalisation, seek treatment at another hospital or substitute other types of care for hospital care. Our estimates of post-strike admissions and mortality will give an indication of whether postponement is a common strategy. In what follows, we also analyse the extent to which surrounding non-university hospitals were affected by the strike.

We use the year 2005 as our benchmark to calculate, for each district, the share of patients who receive treatment in any strike-hit university hospital in that year. We argue that this variable is a good proxy to study strike spillover effects, since the pressure on surrounding non-treated hospitals should be greatest in districts where a high share of patients treated in strike-hit hospitals at baseline. The baseline patient shares is thus used as a regressor in the following regression model:

$$y_{it} = \beta_0 + \beta_1 S_{it} + \mu_i + \nu_t + C_i \cdot t + \epsilon_{it} \quad (3.3)$$

where y_{it} is one of the hospital outcome variables for district i on day t . S_{it} is an interaction between the treatment indicator and the pre-strike share of patients treated in university hospitals. Remaining variables are defined as above.


3.5 Results

3.5.1 Hospital Admissions and Patient Mortality

We first investigate whether the strike affected the number of daily hospital admissions as a consequence of the 2006 physician strike. Figure 3.2 illustrates bi-monthly event study estimates for the relative change in hospital admissions for striking and non-striking hospitals. The absence of a trend for the point estimates in the time periods leading up to the strike suggests that the critical common trend assumption required for a causal interpretation of

our DID estimator is valid. Considering the estimates for the strike period itself, indicated by the vertical dashed lines, we see a significant relative drop in admissions in striking hospitals in May and June 2006 by about 15 cases per day when the labour disruptions were the most intense. Finally, the coefficient estimates once again turn statistically indistinguishable from zero once the strike ends, indicating that striking hospitals experienced a permanent drop in their admission rates as a consequence of the strike.

Figure 3.2: Event Study Results - Hospital Admissions



Notes: The figure shows event study estimates for admissions, controlling for day of week fixed effects and hospital fixed effects. The vertical lines represent the beginning and end of the strike. 95% confidence intervals included. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

Formal DID estimates are reported in Table 3.5, confirming the event study result that the strike significantly decreased hospital admissions. In particular, when including time and hospital fixed effects and hospital-specific linear time trends, the estimates from Columns (1) and (2) suggest a decrease in daily admissions by 19-21 cases on average, or 12 percent. Columns (3) and (4) of the table reports results from applying the covariate adjustment approach described in Section 3.4.2. While this adjustment marginally attenuates the point estimate, it remains large and highly statistically significant.

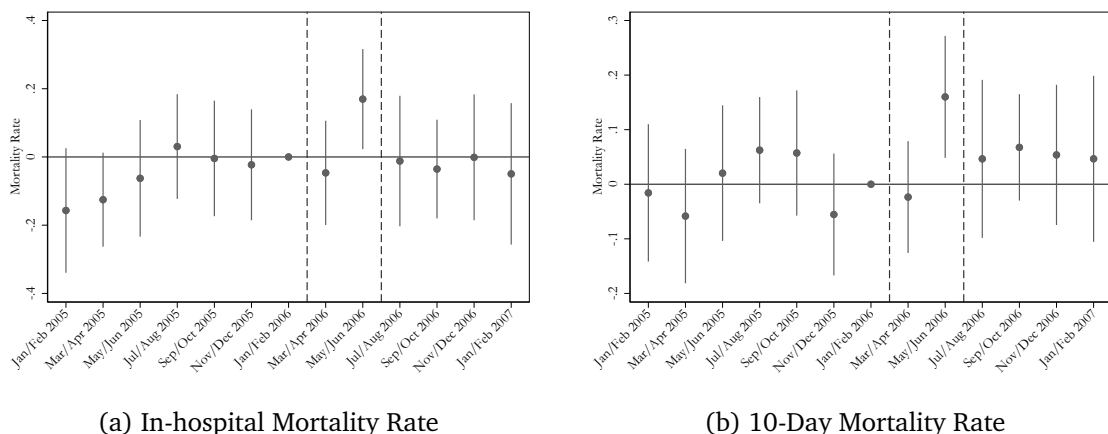
Table 3.5: Difference-in-differences Estimates - Hospital Admissions

	Observations	Mean	(1)	(2)	(3)	(4)
Cases	143,744	158.32	-21.460*** (4.646)	-19.041*** (2.596)	-15.583*** (3.444)	-13.217*** (1.717)
Log Cases	143,744	4.90	-0.122*** (0.019)	-0.118*** (0.014)	-0.082*** (0.015)	-0.079*** (0.010)
Covariate adjustment					✓	✓
Year FE			✓	✓	✓	✓
Hospital FE			✓	✓	✓	✓
Hospital Specific Trends				✓		✓

Notes: Robust standard errors clustered at the hospital level are reported in parenthesis. Significance levels: * 0.10 ** 0.05 *** 0.01. Regressions include day of week fixed effects, week fixed effects, year fixed effects, hospital fixed effects and hospital specific time trends. Regressions are weighted by number of admissions. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

Next, we analyse the consequences of the 2006 strike on in-hospital mortality rates. Figure 3.3 presents event study graphs for the overall mortality rate and 10-day mortality rate. Again we can observe that striking and non-striking hospitals had similar mortality trends in the lead-up to the strike, confirming that our empirical strategy appears sound in this context. We also see a reverse trend in the estimates for the strike months with a relative increase in both in-hospital and 10-day mortality rates for striking hospitals during the intensive strike period in May and June. These estimates return to being close to zero and statistically insignificant at the end of the strike period.

Figure 3.3: Event Study Results - Hospital Mortality



Notes: The figures show event study estimates for mortality rate and 10-day mortality rate, controlling for day of week fixed effects and hospital fixed effects. The vertical lines represent the beginning and end of the strike. 95% confidence intervals included. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

DID estimates corresponding to the event study plots for the mortality outcomes are presented in Columns (1) and (2) of Table 3.6. The point estimates imply that the mortality rate increased by roughly 0.16 percentage points, corresponding to an increase of around nine percent compared to the baseline mean of 1.75 deaths per 100 admissions. The estimates for 10-day mortality are somewhat lower at 0.07 percentage points, or 7.3 percent, but remain significant at the five percent level of statistical significance. Finally, reported estimates from the third row of the table indicate that the number of deaths in hospital drop by 0.07 from a baseline of 2.46; a reduction by 2.8 percent. This estimate is not significantly significant, yet consistent with the estimates for admissions (-12%) and mortality (+9%) since $0.88 \cdot 1.09 = 0.96$.

Columns (3) and (4) of the table reports results from applying the covariate adjusted approach described in Section 3.4.2. The results suggest that, conditional on the selected covariates, the strike had a moderate effect on case-specific mortality rates: the increase is now down to 0.8 percentage points or 4.6 percent. This suggests that a substantial part of the above estimated mortality effect is due to selection of sicker patients into striking hospitals during the strike period. This interpretation is confirmed by the fact that when we estimate the effect of the strike on predicted mortality, we obtain a treatment effect of 0.06 percentage points.


Table 3.6: Difference-in-differences Estimates - Hospital Mortality

	Observations	Mean	Baseline Specification		Covariate Adjusted	
			(1)	(2)	(3)	(4)
Mortality Rate	143,744	1.75	0.159*** (0.043)	0.165*** (0.041)	0.080** (0.037)	0.085** (0.035)
10-Day Mortality Rate	143,744	0.96	0.074** (0.036)	0.073** (0.035)	0.022 (0.030)	0.021 (0.030)
Number of Deaths	143,744	2.46	-0.099 (0.087)	-0.069 (0.056)	-0.088 (0.077)	-0.062 (0.054)
Year FE			✓	✓	✓	✓
Hospital FE			✓	✓	✓	✓
Hospital Specific Trends				✓		✓
Covariate adjustment					✓	✓

Notes: Robust standard errors clustered at the hospital level are reported in parenthesis. Significance levels: * 0.10 ** 0.05 *** 0.01. Regressions include day of week fixed effects, week fixed effects, year fixed effects, hospital fixed effects and hospital specific time trends. Regressions are weighted by number of admissions. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

If there were noteworthy and significant strike effects on mortality, these should also be visible for the district mortality. We therefore present event study estimates for treatment and control districts in Figure 3.4. The graph suggests no significant increase in mortality during the strike period or thereafter.

Figure 3.4: Event Study Results - District Mortality



Notes: The figure shows event study estimates for log mortality, controlling for date fixed effects and district fixed effects. The vertical lines represent the beginning and end of the strike. 95% confidence intervals included. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008b), own calculations.

3.5.2 Care Rationing and Hospital Spillovers

In this section, we explore how the physician strike affected the patient composition and potential spillover effects to other non-striking hospitals. We first study changes in the striking hospitals' patient compositions. Table 3.7 reports estimates for a set of outcomes associated with patient characteristics, including age, sex, rates of Ambulatory Case Sensitive (ACS) and emergency cases, an indicator for urgency (yes/no) and a deferability index, and hospital decisions, including length of stay and rate of surgery. We distinguish between these two sets of outcomes as the latter are more related to hospitals' choices in responding to the strike, whereas the former are exogenous to the hospital.

The estimates presented in the table suggest that the strike led to a significant reduction of patients' age by 0.41 years and significant increases in both emergency rates and the probability that a patient is classified as an

Table 3.7: Difference-in-differences estimates: Patient composition

	Observations	Mean	(1)	(2)
PATIENT COMPOSITION				
Age	143,744	47.89	-0.369** (0.157)	-0.411*** (0.146)
Female	143,744	0.49	0.002 (0.002)	0.002 (0.002)
ACS Rate	143,744	20.30	-0.188 (0.221)	-0.301 (0.183)
Emergency Rate	143,156	28.41	1.558*** (0.349)	1.505*** (0.348)
Urgency Indicator	143,156	0.34	0.012*** (0.003)	0.012*** (0.002)
Deferability Index	143,549	0.12	0.003*** (0.001)	0.003*** (0.001)
HOSPITAL DECISIONS				
Length of Stay	143,744	8.17	0.255*** (0.055)	0.243*** (0.052)
Surgery Rate	143,744	35.95	0.016 (1.409)	-1.629* (0.930)
Year FE			✓	✓
Hospital FE			✓	✓
Hospital Specific Trends				✓

Notes: Robust standard errors clustered at the hospital level are reported in parenthesis. Significance levels: * 0.10 ** 0.05 *** 0.01. Regressions include day of week fixed effects, week fixed effects, year fixed effects, hospital fixed effects and hospital specific time trends. Regressions are weighted by number of admissions. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

urgent case. This is consistent with the fact that emergency care patients in general tend to be younger on average.¹¹ With respect to hospital decisions, we estimate a statistically significant increase in length of stay by 0.24 days. These findings together suggest that patients in relatively poorer health conditions were admitted to striking hospitals during the strike period. We do not find significant changes in the composition of gender or ACS conditions. For surgery rates, we find a considerable, albeit only marginally significant, reduction by about 4.5 percent, suggesting a possible postponement of scheduled surgeries during the strike period.

Next, we split our sample to compare results by emergency status to study the hypothesis in Gruber and Kleiner (2012) where stronger strike effects may

¹¹Figure 3.A.1 in 3.A shows the age distribution of non-emergency versus emergency care patients.

be observed for non-emergency cases if healthier patients delay treatment or go to other health care facilities as a consequence of the strike. Table 3.8 presents separate DID estimates for emergency and non-emergency patients in our sample. Indeed, we see that the decrease in the number of admissions is considerably greater in magnitude for non-emergency cases (15.9%) compared to emergency cases (7.1%) suggesting a relative shift to more urgent cases. In addition, effects on mortality rates are stronger (relative to baseline means) and only significant for non-emergency admissions, and the impact on length of stay is also greater in magnitude for non-emergencies. Overall, the results in Table 3.8 suggest that healthier patients were less likely to receive care at striking hospitals compared to patients with poorer health status and more acute health conditions.

Table 3.8: Difference-in-differences Estimates - Heterogeneity by Emergency Status

	Emergency			Non-Emergency		
	Mean	(1)	(2)	Mean	(1)	(2)
Cases	143.06	-18.516*** (4.197)	-16.631*** (2.294)	164.18	-21.976*** (4.640)	-19.793*** (2.700)
Log Cases	3.48	-0.069*** (0.016)	-0.071*** (0.010)	4.57	-0.159*** (0.023)	-0.155*** (0.020)
Mortality Rate	3.37	0.133 (0.102)	0.143 (0.098)	1.27	0.207*** (0.044)	0.207*** (0.043)
10-day Mortality Rate	2.21	0.071 (0.080)	0.072 (0.079)	0.55	0.080*** (0.026)	0.074*** (0.025)
Number of Deaths	1.19	-0.064 (0.049)	-0.055 (0.038)	1.30	-0.029 (0.046)	-0.015 (0.033)
Length of Stay	9.31	0.241** (0.107)	0.232** (0.092)	8.00	0.368*** (0.061)	0.358*** (0.057)
Surgery Rate	24.10	0.047 (1.220)	-1.256 (0.786)	40.31	0.074 (1.591)	-1.798* (1.063)
Year FE		✓	✓		✓	✓
Hospital FE		✓	✓		✓	✓
Hospital Specific Trends			✓			✓

Notes: Robust standard errors clustered at the hospital level are reported in parenthesis. Significance levels: * 0.10 ** 0.05 *** 0.01. Regressions include day of week fixed effects, week fixed effects, year fixed effects, hospital fixed effects and hospital specific time trends. Regressions are weighted by number of admissions. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

To study spillover effects to other hospitals, we estimate Equation (3.3) for a set of outcomes. Table 3.9 reports the estimated β_1 parameter, interpreted as the average percentage change in the outcome variable for each percentage

point of a hospital's pre-strike patient population that were treated in a striking hospital. The results indicate that the hospitals surrounding the treated university hospitals experience a surge in admissions by about eight percent. However, there is no indication that this increase in admissions worsens patient outcomes as the parameter estimates for mortality rates and length of stay are generally small and non-significant at conventional levels.

Table 3.9: Difference-in-differences Estimates - Spillover Analysis

	Observations	Mean	(1)	(2)
Log Cases	1,041,569	4.25	0.102*** (0.031)	0.078*** (0.020)
Mortality Rate	1,041,569	2.57	-0.118 (0.124)	0.039 (0.094)
10-Day Mortality Rate	1,041,569	1.62	-0.070 (0.090)	-0.014 (0.081)
Length of Stay	1,041,736	8.19	0.483*** (0.143)	0.105 (0.109)
Year FE			✓	✓
Hospital FE			✓	✓
Hospital Specific Trends				✓

Notes: Robust standard errors clustered at the hospital level are reported in parenthesis. Significance levels: * 0.10 ** 0.05 *** 0.01. Regressions include day of week fixed effects, week fixed effects, year fixed effects, hospital fixed effects and hospital specific time trends. Except for log cases regressions are weighted by number of admissions. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

3.5.3 Robustness Checks

In this section, we examine a number of robustness checks to probe the stability of our empirical findings. First, to analyse whether seasonality is a confounding factor, we estimate hospital level placebo regressions for our outcomes dating back the strike to the same period of the year but in 2004. Table 3.10 presents the results showing small and statistically insignificant estimates across the board.

Next, we test the stability of our findings to the choice of control group. To this end, we use all other hospitals in Hamburg and Berlin, who neither participated in the federal nor the municipality strike, as alternative controls.¹²

¹²Berlin did not participate in the strike, because it was not a member of the employers' association and Hamburg reached a collective agreement before the strike began.

Table 3.10: Difference-in-differences estimates: Placebo Regressions for 2004

	Observations	Mean	(1)	(2)
Cases	143,744	158.32	-1.032 (4.081)	-1.311 (4.265)
Log Cases	143,744	4.90	0.009 (0.014)	0.007 (0.015)
Mortality Rate	143,744	1.75	0.021 (0.058)	-0.006 (0.038)
10-day Mortality Rate	143,744	0.96	0.002 (0.037)	-0.015 (0.025)
Number of Deaths	143,744	2.46	0.002 (0.110)	-0.034 (0.101)
Length of Stay	143,744	8.17	-0.021 (0.061)	-0.026 (0.062)
Surgery Rate	143,744	35.95	2.610 (1.853)	2.395 (1.865)
ACS Rate	143,744	20.30	0.056 (0.178)	0.077 (0.180)
Emergency Rate	143,156	28.41	-0.121 (0.204)	-0.106 (0.202)
Urgency Indicator	143,156	0.34	-0.002 (0.002)	-0.001 (0.002)
Deferability Index	143,549	0.12	-0.000 (0.000)	-0.000 (0.000)
Year FE			✓	✓
Hospital FE			✓	✓
Hospital Specific Trends				✓

Notes: Robust standard errors clustered at the hospital level are reported in parenthesis. Significance levels: * 0.10 ** 0.05 *** 0.01. Regressions include day of week fixed effects, week fixed effects, year fixed effects, hospital fixed effects and hospital specific time trends. Except for log cases all regressions are weighted by the number of admissions. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

Note that this alternative control group includes all hospitals in the two federal states, not only university hospitals, and is therefore potentially less comparable with respect to a range of factors, including size and scope of provided services. Nevertheless, Table 3.11 shows that our findings are largely robust to the choice of control group.

Table 3.11: Difference-in-differences Estimates - Alternative Control Group

	Observations	Mean	(1)	(2)
Cases	267,382	161.83	-21.977*** (3.718)	-19.536*** (2.354)
Log Cases	267,382	4.72	-0.138*** (0.016)	-0.125*** (0.013)
Mortality Rate	267,382	2.07	0.270*** (0.045)	0.199*** (0.041)
10-day Mortality Rate	267,382	1.19	0.135*** (0.036)	0.095*** (0.034)
Number of Deaths	267,382	3.07	0.006 (0.090)	-0.035 (0.055)
Length of Stay	267,382	8.28	0.392*** (0.064)	0.261*** (0.053)
Surgery Rate	267,382	37.35	-1.765 (1.607)	-3.004** (1.170)
ACS Rate	84,986	19.25	-0.483*** (0.148)	-0.462*** (0.150)
Emergency Rate	84,648	27.65	1.494*** (0.350)	1.484*** (0.353)
Urgency Indicator	84,648	0.33	0.012*** (0.002)	0.012*** (0.003)
Deferability Index	84,865	0.12	0.003*** (0.001)	0.003*** (0.001)
Year FE			✓	✓
Hospital FE			✓	✓
Hospital Specific Trends				✓

Notes: Robust standard errors clustered at the hospital level are reported in parenthesis. Significance levels: * 0.10 ** 0.05 *** 0.01. Regressions include day of week fixed effects, week fixed effects, year fixed effects, hospital fixed effects and hospital specific time trends. Except for log cases all regressions are weighted by the number of admissions. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

We also test the robustness of our main findings by specifying an alternative treatment indicator, as the one we use in our main analysis does not capture patients that were already hospitalised when the strike started. The alternative treatment indicator is defined as the share of patients on a specific day that were affected by the strike at any time during their hospital stay. The reported estimates from Table 3.12 show that our results are robust to using this alternative definition of treatment group.

Table 3.12: Difference-in-differences Estimates - Alternative Treatment Indicator

	Observations	Mean	(1)	(2)
Cases	143,744	158.32	-20.065*** 4.984	-17.453*** 2.645
Log Cases	143,744	4.90	-0.107*** 0.019	-0.103*** 0.014
Mortality Rate	143,744	1.75	0.128*** 0.044	0.134*** 0.042
10-day Mortality Rate	143,744	0.96	0.055 0.036	0.054 0.035
Number of Death	143,744	2.46	-0.097 0.095	-0.064 0.062
Length of Stay	143,744	8.17	0.226*** 0.058	0.217*** 0.054
Surgery Rate	143,744	35.95	-0.473 1.510	-1.856* 0.945
ACS Rate	143,744	20.30	0.009 0.233	-0.134 0.193
Emergency Rate	143,156	28.41	1.364*** 0.334	1.312*** 0.329
Urgency Indicator	143,156	0.34	0.010*** 0.002	0.010*** 0.002
Deferability Index	143,549	0.12	0.003*** 0.001	0.003*** 0.001
Year FE			✓	✓
Hospital FE			✓	✓
Hospital Specific Trends				✓

Notes: Robust standard errors clustered at the hospital level are reported in parenthesis. Significance levels: * 0.10 ** 0.05 *** 0.01. Regressions include day of week fixed effects, week fixed effects, year fixed effects, hospital fixed effects and hospital specific time trends. Except for log cases all regressions are weighted by the number of admissions. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

Finally, we also test for linear pre-trends by estimating the following regression model on the pre-intervention sample:

$$y = \beta(\text{trend} \times \text{treated}) + \gamma\text{treated} + \delta\text{trend} + \epsilon, \quad (3.4)$$

where *trend* is a non-parametric time trend estimated for each *year* × *month* cell in the sample, and *treated* is a treatment indicator equal to one for striking hospitals and zero for non-striking hospitals. If the estimate β is significantly different from 0, we conclude the treated and control hospitals had significantly different trends leading up to the strike; thus rejecting the common trend assumption required for consistency of the DID estimator. The

reported $\hat{\beta}$ in Table 3.13 indicate that the common trend assumption cannot be rejected at any conventional levels of statistical significance.¹³ In addition, it should also be noted that our main specification includes hospital-specific time trends which makes differential pre-trends less likely to be a problem.

Table 3.13: Pre-trend Test

	(1)	(2)	(3)	(4)
	Log Cases	Mortality Rate	10-Day Mortality Rate	Cases
Trend × Treated	-0.0000262	-0.0000480	-0.0000434	-0.00972
	(-0.81)	(-1.07)	(-1.65)	(-0.92)

Notes: Robust standard errors clustered at the hospital level are reported in parenthesis. Significance levels: * 0.10 ** 0.05 *** 0.01. Except for log cases all regressions are weighted by the number of admissions. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

3.6 Conclusion

Strikes in essential industries, such as transport, healthcare and protection, have become more common in recent times due to budget pressures and worker dissatisfaction with rationalisation and retrenchment policies in the public sector. In the healthcare sector, longer working hours, less job security and higher workloads has increased the intensity of industrial action, which causes social and individual costs in the form of lower access and quality of care. Given the high-stake nature of the occupation, it is an important task to study the impact of strikes in the healthcare sector; in particular among specialists including physicians.

We explore the short-term impacts of a nationwide physician strike in German university hospitals in 2006 on patient mortality risk. We compare changes in outcomes over time in hospitals that were subject to striking physicians to other hospitals in a difference-differences empirical design. Our results show that hospital admissions in striking hospitals dropped by an estimated 12 percent during the strike period compared to non-striking hospitals. Moreover, we find that in-hospital mortality rates increased by nine

¹³The recent literature recommends caution regarding pre-trend tests as their power may be low (Roth et al., 2022). Nevertheless, the test does provide additional evidence concerning the plausibility of the common time trend assumption.

percent over the same period. However, the effect on mortality is attenuated after adjusting for patient case-mix, suggesting that the effect is also due to a change in the patient composition during the strike period. This result is indicative of a triaging strategy where more urgent or frailer cases gain precedence over more healthy patients. Finally, we find no indications of neither temporal (post-strike period) or spatial (other nearby hospitals) spillover admission effects from which we conclude that the drop in admissions in striking hospitals were most likely permanent.

Our results contrast sharply with the related strike literature in Economics. For example, Gruber and Kleiner (2012), find a 26 percent decrease in admissions during the nurse strikes and an increase of 18 percent in in-hospital mortality. We interpret this glaring difference as an indication that strikes in healthcare (and elsewhere) can have very different impacts on organisational efficacy depending on the extent and range of services affected by the disruption. One drawback of the analysis in this paper is that death is a rare and serious outcome (Cunningham et al., 2008). It is still possible, and even likely, that other quality or patient care indicators, such as readmission rates or patient satisfaction, were affected. Furthermore, there might be other economic consequences of strikes, such as financial losses for the hospitals or an increased workload for non-striking physicians that we do not explore further in this paper. Establishing the total impact of strikes on a fuller set of social, economic, and health indicators should therefore constitute an important avenue for further research.

Bibliography

- Aiken, L. H., Clarke, S. P., Sloane, D. M., Sochalski, J., and Silber, J. H. (2002). Hospital nurse staffing and patient mortality, nurse burnout, and job dissatisfaction. *Jama*, 288(16):1987–1993.
- Avdic, D. (2016). Improving efficiency or impairing access? health care consolidation and quality of care: Evidence from emergency hospital closures in sweden. *Journal of health economics*, 48:44–60.
- Belloni, A., Chernozhukov, V., and Hansen, C. (2014). Inference on treatment effects after selection among high-dimensional controls. *The Review of Economic Studies*, 81(2):608–650.
- Blümel, M. and Busse, R. (2015). The german health care system, 2015. *International profiles of health care systems*, pages 69–76.
- Bundesministerium für Gesundheit (2019). Daten des gesundheitswesens 2019. *Bundesministerium für Gesundheit*. <https://www.bundesgesundheitsministerium.de/themen/krankenversicherung/zahlen-und-fakten-zur-krankenversicherung/kennzahlen-daten-bekanntmachungen.html> Retrieved on April 21, 2020.
- Costa, E. (2019). License to kill? the impact of hospital strikes. *The Impact of Hospital Strikes (May 20, 2019)*.
- Cunningham, S. A., Mitchell, K., Narayan, K. V., and Yusuf, S. (2008). Doctors' strikes and mortality: a review. *Social Science & Medicine*, 67(11):1784–1788.
- de Cruppé, W. and Geraedts, M. (2017). Hospital choice in germany from the patient's perspective: a cross-sectional study. *BMC health services research*, 17(1):720.
- Dribbusch, H. (2018). Das einfache, das so schwer zu zählen ist: Probleme der streikstatistik in der bundesrepublik deutschland. *Industrielle Beziehungen*, 25(3).

- FDZ der Statistischen Ämter des Bundes und der Länder (2000-2008a). Krankenhausstatistik der Jahre 2000-2008.
- FDZ der Statistischen Ämter des Bundes und der Länder (2000-2008b). Todesursachenstatistik der Jahre 2000-2008.
- Federal Statistical Office (2018). Grunddaten der Krankenhäuser 2017, Fachserie 12 Reihe 6.1.1. *Statistisches Bundesamt (Destatis)*.
- Friedman, W. and Keats, A. (2019). Disruptions to health care quality and early child health outcomes: Evidence from health worker strikes in Kenya. mimeo.
- Goepfert, A. and Conrad, C. B. (2013). *Unternehmen Krankenhaus*. Georg Thieme Verlag.
- Greef, S. (2012). *Die Transformation des Marburger Bundes: Vom Berufsverband zur Berufsgewerkschaft*. Springer-Verlag.
- Gruber, J. and Kleiner, S. A. (2012). Do strikes kill? Evidence from New York State. *American Economic Journal: Economic Policy*, 4(1):127–57.
- Hirani, J., Sievertsen, H. H., and Wüst, M. (2019). Beyond treatment exposure: The timing of early interventions and children's health. mimeo.
- International Labour Office (2006). Freedom of association and others. Digest of decisions and principles of the Freedom of Association Committee of the Governing Body of the ILO. *International Labour Office, Geneva*.
- KHStatV (1990). Krankenhausstatistik-Verordnung vom 10. April 1990 (BGBI. I S. 730), die zuletzt durch Artikel 1 der Verordnung vom 10. Juli 2017 (BGBI. I S. 2300) geändert worden ist. *Bundesamt für Justiz*.
- Krämer, J., Schreyögg, J., and Busse, R. (2019). Classification of hospital admissions into emergency and elective care: A machine learning approach. *Health Care Management Science*, 22(1):85–105.
- Kronborg, H., Sievertsen, H. H., and Wüst, M. (2016). Care around birth, infant and mother health and maternal health investments—evidence from a nurse strike. *Social Science & Medicine*, 150:201–211.

- Lin, H. (2014). Revisiting the relationship between nurse staffing and quality of care in nursing homes: An instrumental variables approach. *Journal of health economics*, 37:13–24.
- Marburger Bund Bundesverband (2012). Marburger bund bundesverband chronologie 2005-2007. Technical report, Marburger Bund Bundesverband.
- Martens, H. (2008). Primäre arbeitspolitik und gewerkschaften im gesundheitswesen: Der ärztestreik 2006 als beispiel primärer arbeitspolitik in zeiten tiefgreifender gesellschaftlicher umbrüche. Technical report, Arbeitspapier 143, Hans-Böckler Stiftung.
- Metcalf, D., Chowdhury, R., and Salim, A. (2015). What are the consequences when doctors strike. *BMJ*, 351:h6231.
- Piérard, E. (2014). The effect of physician supply on health status: Canadian evidence. *Health Policy*, 118(1):56–65.
- Roth, J., Sant’Anna, P. H., Bilinski, A., and Poe, J. (2022). What’s trending in difference-in-differences? a synthesis of the recent econometrics literature. *arXiv preprint arXiv:2201.01194*.
- Stoye, G. and Warner, M. (2023). The effects of doctor strikes on patient outcomes: Evidence from the english nhs. *Journal of Economic Behavior & Organization*, 212:689–707.
- Sundmacher, L., Fischbach, D., Schuettig, W., Naumann, C., Augustin, U., and Faisst, C. (2015). Which hospitalisations are ambulatory care-sensitive, to what degree, and how could the rates be reduced? results of a group consensus study in germany. *Health Policy*, 119(11):1415–1423.

Appendix

3.A Appendix Figures and Tables

Table 3.A.1: Physicians and Health Care Workers in Hospitals


Year	All Hospitals			University Hospitals		
	Total	Physicians	Other Medical	Total	Physicians	Other Medical
2000	1,100,471	122,062	978,409	172,867	24,398	148,469
2001	1,101,356	123,819	977,537	173,114	24,758	148,356
2002	1,112,421	126,047	986,374	174,850	25,084	149,766
2003	1,096,420	128,853	967,567	173,091	25,154	147,937
2004	1,071,846	129,817	942,029	168,980	25,171	143,809
2005	1,063,154	131,115	932,039	170,160	25,435	144,725
2006	1,064,377	133,649	930,728	171,895	25,781	146,114
2007	1,067,287	136,267	931,020	172,782	26,241	146,541
2008	1,078,212	139,294	938,918	173,182	26,488	146,694
2009	1,096,520	143,967	952,553	177,535	27,632	149,903
2010	1,112,959	148,696	964,263	181,954	28,443	153,511

Notes: Information System of the Federal Health Monitoring.

Table 3.A.2: Unweighted Mortality Regressions

	Observations	Mean	Estimate
Mortality Rate	143,744	2.15	0.117* (0.065)
10 Day Mortality Rate	143,744	1.23	0.084 (0.050)
Year FE			✓
Hospital FE			✓
Hospital Specific Trends			✓

Notes: Robust standard errors clustered at the hospital level are reported in parenthesis. Significance levels: * 0.10 ** 0.05 *** 0.01. Regressions include day of week fixed effects, week fixed effects, year fixed effects, hospital fixed effects and hospital specific time trends. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

Figure 3.A.1: Age Distribution by Emergency Status

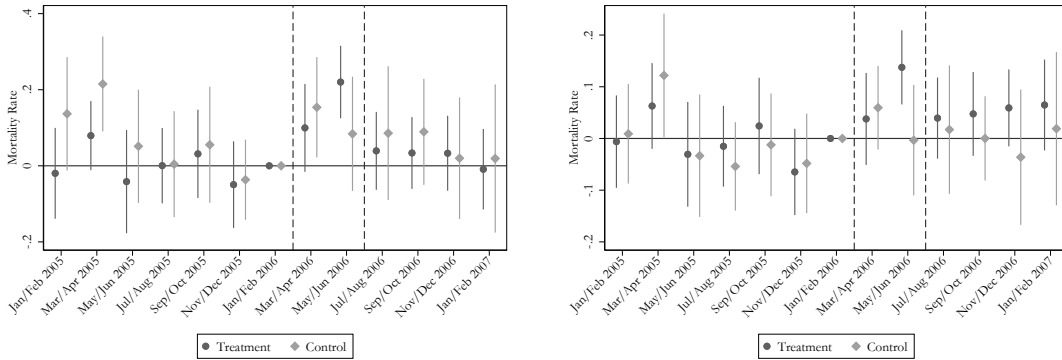
Notes: The figure shows the age distribution by emergency status. Based on individual level hospital data. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

Table 3.A.3: Admissions by Cause

	Observations	Mean	(1)	(2)
Cardio	143,744	20.75	-2.465*** (0.816)	-2.323*** (0.464)
Respiratory	143,744	7.17	-1.093*** (0.320)	-0.914*** (0.221)
Infectious	143,744	3.40	-0.287** (0.139)	-0.183** (0.082)
Metabolic	143,744	4.68	-0.636*** (0.125)	-0.544*** (0.089)
Neoplastic	143,744	33.54	-3.622*** (0.846)	-2.844*** (0.511)
Birth	143,744	7.96	-0.026 (0.146)	-0.036 (0.139)
Medical Complication	143,744	0.62	-0.114** (0.056)	-0.111** (0.045)
Year FE			✓	✓
Hospital FE			✓	✓
Hospital Specific Trends				✓

Notes: This table shows admissions by admission cause. Robust standard errors clustered at the hospital level are reported in parenthesis. Significance levels: * 0.10 ** 0.05 *** 0.01. Regressions include day of week fixed effects, week fixed effects, year fixed effects, hospital fixed effects and hospital specific time trends. Except for log cases regressions are weighted by number of admissions. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

Figure 3.A.2: Descriptives Mortality Rate




(a) In-hospital Mortality Rate

(b) 10-Day Mortality Rate

Notes: The figures show means for treatment and control hospitals. The vertical lines represent the beginning and end of the strike. 95% confidence intervals included. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

Figure 3.A.3: Descriptives Cases

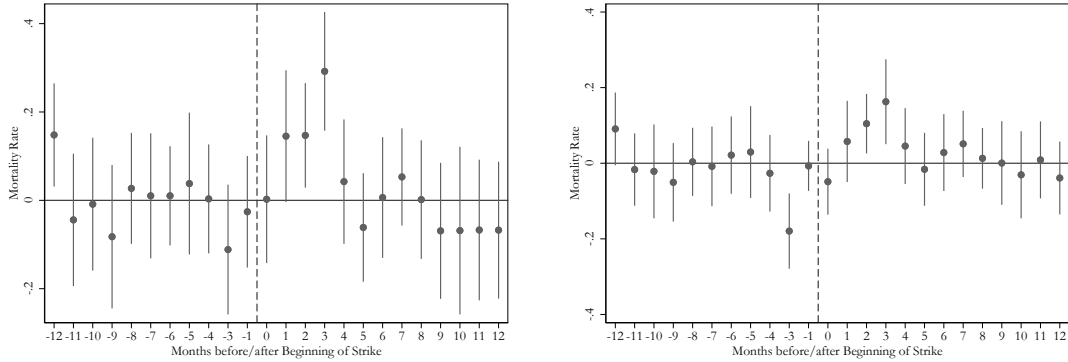


(a) Cases

(b) Log Cases

Notes: The figures show means for treatment and control hospitals. The vertical lines represent the beginning and end of the strike. 95% confidence intervals included. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

Figure 3.A.4: Alternative Estimator - Mortality Rate




(a) In-hospital Mortality Rate

(b) 10-Day Mortality Rate

Notes: The figures show event study estimates for mortality rate and 10-day mortality rate, controlling for day of week fixed effects and hospital fixed effects using the estimator of ?. Reference period $t - 2$. The vertical lines represent the beginning and end of the strike. 95% confidence intervals included. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

Figure 3.A.5: Alternative Estimator - Cases



(a) Cases

(b) Log Cases

Notes: The figures show event study estimates for cases and log cases, controlling for day of week fixed effects and hospital fixed effects using the estimator of ?. Reference period $t - 2$. The vertical lines represent the beginning and end of the strike. 95% confidence intervals included. Source: FDZ der Statistischen Ämter des Bundes und der Länder (2008a), own calculations.

3.B Variable Definitions

- **Treatment Variable:** Dummy variable taking on the value 1 for admissions to a striking hospital during the strike period.
- **Female:** Dummy variable taking on the value 1 for females.
- **Age:** Age in years, continuous variable.
- **Married:** Dummy variable taking on the value 1 for married individuals.
- **Cases:** Number of admissions per day per hospital.
- **Log Cases:** Logarithmised number of admissions per day per hospital.
- **Length of Stay:** Length of hospital stay in days.
- **Surgery Rate:** Surgeries per 100 admissions.
- **Mortality Rate:** Deaths per 100 admissions.
- **10 Day Mortality Rate:** Deaths within 10 days after admission per 100 admissions.
- **ACS Rate:** Ambulatory care sensitive admissions based on (Sundmacher et al., 2015) per 100 admissions.
- **Urgency Indicator:** Urgency indicator based on Krämer et al. (2019), continuous variable taking on values between 0 and 1.
- **Emergency Rate:** Admissions indicated as emergency based on Krämer et al. (2019) per 100 admissions. Derived from urgency indicator. Admission is defined as an emergency admission if urgency indicator is above 0.5.

Concluding Remarks

In this thesis, I investigate the effect of three different events: job losses, moves and strikes – by using large administrative data.

First, I provide evidence that job loss has long-lasting effect, even affecting workers' retirement decision. Using German administrative data spanning from 1975 to 2021, I exploit firm closures as a natural experiment to compare the retirement behavior of workers who experienced a job displacement with similar workers who did not.

My analysis reveals that displaced workers postpone retirement in response to job loss. Nonetheless, despite their adjustments, displaced workers continue to face substantial declines in their estimated pension benefits. The lifetime costs associated with job displacement are substantial, with displaced workers experiencing reductions in their present discounted value of income amounting to approximately 16%. These findings shed light on the interplay between job loss and retirement choices, emphasizing the necessity for comprehensive policy considerations in mitigating the long-term impacts of job loss.

In chapter 2, we show that the gender gap in earnings after relocation cannot be explained only by gender differences in earnings or potential earnings but rather a norm prioritizing men's careers.

Employing an event-study approach, we first show that relocation yields more substantial earnings gains for men compared to women, both in Germany and Sweden. To explore the underlying drivers of these gendered patterns, we construct a model of household decision-making that allows households to potentially put less weight on the income earned by the woman relative to the man. We test this model using a subset of couples with similar potential earnings. Our analysis reveals that households in both countries place less weight on income earned by the woman compared to the man, particularly in Germany. This underscores the interplay between gender norms, economic incentives, and household decision-making within the context of couples navigating career choices and location decisions.

Finally, the last chapter shows that physician strikes have large negative health effects. Using an extensive dataset covering all hospital admissions in

Germany and digitized records of strike involvement, we apply a difference-in-differences model to show that the strike leads to a significant reduction in hospital admissions and an increase in mortality. However, the effects on in-hospital mortality predominantly stem from patient selection.

We show that emergency cases and patients with greater fragility⁹ were more likely to remain hospitalized during the strike period. Consequently, our results suggest that, contrary to expectations, short-term disruptions in healthcare access may not significantly harm healthcare quality, provided that care allocation, based on patient severity, is effectively executed. This underscores the importance of adaptability and prioritization within healthcare systems during periods of service interruptions.

Erklärung

gemäß §10 Abs. 6 der Promotionsordnung der Mercator School of Management, Fakultät für Betriebswirtschaftslehre der Universität Duisburg-Essen, vom 11. Juni 2012.

Hiermit versichere ich, dass ich die vorliegende Dissertation selbständig und ohne unerlaubte Hilfe angefertigt und andere als die in der Dissertation angegebenen Hilfsmittel nicht benutzt habe. Alle Stellen, die wörtlich oder sinngemäß aus anderen Schriften entnommen sind, habe ich als solche kenntlich gemacht.

Duisburg, 16. Oktober 2023

Lea Nassal

DuEPublico

Duisburg-Essen Publications online

UNIVERSITÄT
DUISBURG
ESSEN

Offen im Denken

ub | universitäts
bibliothek

Diese Dissertation wird via DuEPublico, dem Dokumenten- und Publikationsserver der Universität Duisburg-Essen, zur Verfügung gestellt und liegt auch als Print-Version vor.

DOI: 10.17185/duepublico/81898

URN: urn:nbn:de:hbz:465-20240503-120712-6

Alle Rechte vorbehalten.