How Do Workers Adjust to Labor Market Shocks?
Essays in Empirical Labor Economics

ISBN: 978 90 3610 590 3

Cover design: Lieke van den Berge
Print: Haveka, Alblasserdam

How Do Workers Adjust to Labor Market Shocks?
Essays in Empirical Labor Economics

Hoe passen werknemers zich aan na een schok op de arbeidsmarkt?
Essays in empirische arbeidseconomie

Thesis

to obtain the degree of Doctor from the
Erasmus University Rotterdam
by command of the
rector magnificus

Prof.dr. R.C.M.E. Engels

and in accordance with the decision of the Doctorate Board.
The public defence shall be held on

Friday 29 November 2019 at 11:30 hrs
by

Adrianus Willem van den Berge

born in Dordrecht

Erasmus University Rotterdam

Doctoral Committee:

Promotor:	Prof.dr.ir. J.C. van Ours
Other members:	Prof.dr. A.J. Dur Prof.dr. B. van der Klaauw Dr. A.C. Gielen

Acknowledgements

Throughout my work on this dissertation I have received a lot of support and help. It would not have been possible to write my thesis without all the people around me, both those who contributed directly to the research included here and those who have been important to me in other areas of life. Here I would like to take the opportunity to thank a few people.

First of all, I would like to thank my supervisor, prof. dr. Jan van Ours for coaching and guiding me throughout the process. Even though we ultimately did not work on a project together, I found your comments and support very helpful.

Second, I would like to thank my co-authors on the projects included in this thesis: James Bessen, Maarten Goos, Egbert Jongen, Anna Salomons, Lisette Swart and Karen van der Wiel. Without you it would not have been possible to complete this work, and it certainly would not have been as enjoyable. Lisette, thank you for always being able to ask me the right (and sometimes difficult) questions and for being a great sparring partner. Anna, your energy and speed of thinking are sometimes overwhelming, but always inspiring. Thank you for taking up the role of my academic mentor and for your inane jokes. Maarten, thank you for remaining level headed whenever we go off the rails and keeping the bigger picture in mind. Jim, thank you for pushing us to work on the important questions. Egbert, I appreciate that you always somehow find the time in your schedule to discuss research or personal things. I hope some of your optimism in planning and publishing will rub off on me. Karen, your coaching helped me a lot when starting out. I have learned a lot from all of you and I hope we will continue our collaboration in the future.

I would also like to thank the members of the reading committee, Robert Dur, Bas van der Klaauw and Anne Gielen, for carefully reading the manuscript and providing valuable comments.

I am indebted to my employer, the CPB, for giving me the opportunity to pursue my PhD. In particular I would like to thank Bas ter Weel, Daniel van Vuuren, Marloes de Graaf-Zijl, Debby Lanser, Egbert Jongen and Laura van Geest, for stimulating and supporting me. I have received generous amounts of time and opportunities to pursue projects that were relevant for my PhD.

This work could not have been completed without the data access provided by Statistics Netherlands. Thanks to all the people behind the scenes answering my questions on the data. And to the CBS for allowing me to access the data for the final two chapters on my own computer instead of in a lonely, dark, cold or hot room (depending on the season).

Some of the time working on the projects in this thesis, in particular the first chapter, was spent at Boston University in the Spring of 2018 and 2019. This was a wonderful learning experience. I would like to thank Jim for hosting me and Jeroen Hinloopen and Debby Lanser for generously allowing me to spend several months in Boston. Of course I thank the two fellow members of Team Dial-a-Llama for making my stay there both intellectually stimulating and a lot of fun. I really enjoyed our lunches, dinners and hiking and kayaking trips. I hope we will continue our research visits in the future.

My PhD was somewhat atypical in that I worked at CPB instead of at university. Here, I was surrounded by a great bunch of smart and committed people who not just gave me an inspiring place to work, but also a really nice and nerdy social environment. We often joke that social skills are not our strong suit, because if they were, all of us would be working at McKinsey. Nevertheless, many of my colleagues have now become close friends. At the risk of forgetting someone, I would like to single out a few people. First of all, Lisette, thank you for all the happy, interesting and sad talks we have had over the years, for your ability to make me rethink things, the support you gave me throughout, and of course for being my paranymph. Benedikt, Bram, Bas, Iris, Sander, Remco, Minke, Marente and Rutger thank you for all the fun times we have had during drinks, karaoke nights, Christmas parties, Sinterklaas celebrations, skiing holidays, study trips, board game nights and all the other activities we either organized or enjoyed. You have certainly provided a lot of happy distractions from writing code and papers. I hope we will continue to have great times together.

Finally, I would like to thank my mother, sister and Martijn for always supporting me no matter what and my father for having always supported me in my choices and life. The last ten years have been far from easy. I am proud to see how strong we have been as a family and I am glad that we have also had many happy moments to share throughout, in particular the arrival of my beautiful niece Chey. Additionally, thank you Lieke for being my paranymph and for designing the cover.

Contents

Acknowledgements **i**

1 Introduction and main findings **1**

2 Automatic Reaction: What Happens to Workers at Firms that Automate? **7**

2.1 Introduction . . . 7
2.2 Data . . . 10
2.3 Empirical approach . . . 13
2.3.1 Defining automation cost spikes . . . 13
2.3.2 Summary statistics on automation cost spikes . . . 16
2.3.3 How do automating firms differ? . . . 17
2.3.4 Empirical design . . . 17
2.4 The impact of automation on incumbent workers . . . 24
2.4.1 Impacts on wage income, firm separation, and non-employment 25
2.4.2 Where do automation-affected workers go? . . . 31
2.4.3 Robustness checks . . . 36
2.4.4 Effect heterogeneity . . . 40
2.5 Comparison to computerization . . . 46
2.5.1 Summary statistics . . . 47
2.5.2 Automation versus computerization . . . 48
2.6 Conclusion . . . 52
2.7 Appendix . . . 54
2.7.1 Sample construction . . . 54
2.7.2 Additional descriptives . . . 55
2.7.3 Additional robustness checks . . . 62

3 Bad Start, Bad Match? The Early Career Effects of Graduating in a Recession for Vocational and Academic Graduates 69
3.1 Introduction . . . 69
3.2 Empirical strategy and data . . . 73
3.2.1 Economic conditions measure and model . . . 74
3.2.2 Selection bias . . . 76
3.2.3 The Dutch higher education system . . . 77
3.2.4 Data sources and sample . . . 79
3.2.5 Constructing the change in field-specific employment . . . 80
3.2.6 Constructing other dependent variables . . . 81
3.2.7 Descriptive statistics . . . 83
3.3 Wage and employment effects of graduating during a recession . . . 84
3.4 Mechanisms of recovery . . . 89
3.4.1 Firm and sector mobility . . . 91
3.4.2 Match quality and the job ladder . . . 92
3.4.3 Returns to job mobility . . . 94
3.5 Conclusion . . . 97
A Appendix . . . 98
A.1 Selection bias . . . 98
A.2 Other rubustness checks . . . 102
A.3 Estimation results for figures in main text . . . 108
A.4 Descriptives on each field of study . . . 112
A.5 Descriptives on firm rank, job mobility and match quality . . . 114

4 Do Parents Work More When Children Start School? Evidence from the Netherlands 117
4.1 Introduction . . . 117
4.2 Theoretical Framework . . . 121
4.3 Magnitude of the treatment . . . 124
4.3.1 Institutional setting . . . 125
4.3.2 Magnitude of the change in terms of time . . . 125
4.3.3 Magnitude of the change in terms of money . . . 126
4.4 Empirical Approach . . . 128
4.4.1 Defining treatment and control group . . . 129
4.4.2 A difference-in-differences design . . . 129
4.4.3 Data and descriptive statistics . . . 131

4.5 Results . . . 134
4.5.1 Main results . . . 134
4.5.2 Heterogeneity . . . 138
4.5.3 Robustness analyses . . . 140
4.6 Conclusion and discussion . . . 142
A Appendix . . . 145
A.1 Description of maternal budget constraint . . . 145
A.2 Additional results . . . 147

5 Using Tax Deductions to Promote Lifelong Learning: Real and Shifting Responses 153
5.1 Introduction . . . 153
5.2 Institutional setup . . . 157
5.3 Theoretical framework . . . 159
5.4 Empirical methodology . . . 160
5.4.1 Singles: regression kink design . . . 161
5.4.2 Couples: donut regression discontinuity design . . . 162
5.5 Data . . . 163
5.6 Results . . . 168
5.6.1 Singles . . . 168
5.6.2 Couples . . . 174
5.6.3 Discussion . . . 177
5.7 Conclusion . . . 180
A Appendix . . . 182

6 Summary and conclusions 201

Samenvatting (Dutch Summary) 203

About the author 207

Portfolio 209

Bibliography 211

List of Tables

2.1 Automation cost share distribution 12
2.2 Automation costs by sector 13
2.3 Automation costs by firm size class 13
2.4 Firm-level automation spike frequency 18
2.5 Automation spike frequency by sector 18
2.6 Automation expenditures by firm type 18
2.7 Relative wage income effects by incumbents' characteristics 43
2.8 Relative wage income effects by incumbents' wage quartile 45
2.9 Computer and automation cost share distributions 49
2.10 Automation costs and computer investments by sector 49
2.11 Automation costs and computer investments by firm size 49
2.12 Automation costs and computer investments by firm size 50
2.13 Number of treatment and control events at the firm level by calendar year 56
2.14 Automation costs by firm size class after removing outliers 58
2.15 Correlates of a firm ever having an automation spike 59
2.16 Brier skill scores for predicting automation spikes 59
2.17 Descriptives for all workers 60
2.18 Descriptives on matched worker samples 61
3.1 Descriptive statistics. 86
3.2 Fixed effects estimates of the wage return to firm and sector mobility. . 96
A1 OLS estimates of the effect of the decline in employment at graduation on the probability of obtaining an additional degree. 99
A2 OLS estimates of the effect of the change in employment at graduation on the composition of the graduation cohort. 100
A3 First stage estimates for IV estimates. 101

A4 IV estimates of the effect of two different indicators for economic conditions at graduation on ln(daily wage). The first measure uses the national change in employment. The second uses the sector-specific change in added value as input in calculating field-specific economic conditions. 104
A5 OLS estimates of the effect of the field-specific decline in employment on ln(daily wage). 105
A6 IV estimates of the effect of the decline in employment at graduation on ln(daily wage) for those who remained within their initial track from secondary education and for those who are observed as employed for each year since graduation. 106
A7 IV estimates using different indicators for match quality. 107
A8 Parameter estimates from IV regressions for Figure 3.5. 109
A9 Parameter estimates from IV regressions for Figure 3.6. 110
A10 Parameter estimates from IV regressions for Figure 3.7. 111
A11 Descriptive statistics on the change in employment at graduation for each field of study. 113
A12 Descriptive statistics on job mobility. 115
A13 Descriptive statistics on firm rank. 116
4.1 Magnitude of the shock in terms of money and in terms of time 127
4.2 Descriptives of demographics for mothers and fathers in treatment and control group weighted by matching weights. 133
4.3 Heterogeneity by hours worked and wage quartile. 140
4.4 Treatment effect estimates on hours per week, hours worked if employed and the probability to work when including self-employed and relative to a control group consisting of parents of children with a second-youngest child aged 3–6. 143
A1 Magnitude of the shock in terms of money and in terms of time 148
A2 Heterogeneity by marital status, ethnicity and number of children. . . . 149
A3 Descriptives of demographics for mothers and fathers in treatment and control group weighted by matching weights for the sample including self-employed. 150
A4 Descriptives of demographics for mothers and fathers in treatment and control group weighted by matching weights for the sample using a control group of parents whose second-youngest child is between three and six years old. 151

5.1 Marginal tax rates and income brackets: 2006–2013 158
5.2 Shifting of lifelong learning expenditures in couples (in %) 165
5.3 The distribution of the use of the deductible 165
5.4 Descriptive statistics: singles 167
5.5 Descriptive statistics: couples 167
5.6 Treatment effect estimates for singles on the probability to use the deductible and the deducted amount, for different bandwidths around the kink 171
5.7 Treatment effect by demographic characteristics for singles at kink 2 . 173
5.8 Treatment effect estimates for singles on the probability to participate in and pay for training: Labour Force Survey 178
5.9 Enrollment in publicly funded education (%) 179
5.10 Frequency of words reported in tax filings 179
A1 Treatment effect estimates for primary earners on the probability to use the deductible and the deducted amount, for different widths of the donut hole 183
A2 Treatment effect estimates secondary earners on the probability to use the deductible and the deducted amount, for different widths of the donut hole 184
A3 Full estimation results for the preferred specification for singles 185
A4 Treatment effect estimates for singles: standard errors 'clustered' at the individual level 185

List of Figures

2.1 Firm-level automation cost shares over time . . . 14
2.2 Firm-level automation cost per worker over time . . . 14
2.3 Automation cost share spikes . . . 19
2.4 Automation cost share spikes for treated firms . . . 19
2.5 Automation cost level per worker for treated firms . . . 20
2.6 Log employment for firms with and without automation events . . . 21
2.7 Log wage bill for firms with and without automation events . . . 22
2.8 Annual real wage income, relative to $t = -1$ . . . 26
2.9 Firm separation hazard . . . 29
2.10 Annual number of days in non-employment . . . 29
2.11 Log daily wages . . . 31
2.12 Probability of switching industries . . . 33
2.13 Annual real benefit income . . . 34
2.14 Cumulative probability of entering self-employment or early retirement 36
2.15 A randomization test for relative wage income estimates . . . 39
2.16 Robustness to removing other firm-level events . . . 39
2.17 Relative annual wage income effects for incumbents versus recent hires 46
2.18 Firm-level computer investment per worker over time . . . 50
2.19 Computer investment per worker for treated firms . . . 51
2.20 Relative annual wage income effects of automation and computerization 52
2.21 Automation cost share spikes for treated firms, balanced sample . . . 57
2.22 Automation cost level per worker for treated firms, balanced sample . . 58
2.23 Annual real wage income in levels . . . 62
2.24 Additional randomization tests . . . 64
2.25 Robustness to changes in model specification . . . 65
2.26 Robustness to removing other firm events . . . 66
2.27 Robustness to changes in spike definition . . . 67
3.1 The share of workers in each aggregated sector for the 5 largest fields of study at each level. . . . 81

3.2 Percentage change in employment for aggregated sectors. 82
3.3 The change in employment over time for the 5 largest fields of study at each level. 84
3.4 Wage-experience profiles for graduates from universities of applied science and university graduates. 85
3.5 Estimated effects of the decline in field-specific employment on wage and employment status. 88
3.6 Estimated effects of the decline in field-specific employment on mean firm wage and job mobility. 93
3.7 Estimated effects of the decline in field-specific employment on match quality and firm rank. 94
A1 Marginal effects of the decline in field-specific employment on ln(daily wage) using a quadratic specification for the decline in field-specific employment. 102
A2 Estimated effects of the decline in field-specific employment on ln(daily wage) for workers graduating in a downturn or an upturn and graduating in the Great Recession or before. 103
4.1 Hours worked by mothers before and after the youngest child in a household starts going to school 119
4.2 The mother's budget constraint. 123
4.3 Descriptive statistics on employment for mothers and for fathers 135
4.4 Main estimates for mothers and fathers 137
5.1 Effective costs of 2,500 euro lifelong learning expenditures 162
5.2 Density around the kink, probability to use the deductible and the deducted amount: singles 169
5.3 Own use of the deductible and own amount: primary earners 175
A1 Declared deductible for primary and secondary earners 182
A2 Declared deducted amount for primary and secondary earners 186
A3 Own use of the deductible and own amount for secondary earners . . . 187
A4 Average (statutory) marginal tax rates in subsequent years for the sample around the kink in 2006 188
A5 RKD plots for control variables for singles 189
A5 RKD plots for control variables for singles (cont.) 190
A6 RD plots for control variables for the primary earner in a couple 191
A6 RD plots for control variables for the primary earner in a couple (cont.) 192
A7 RD plots for control variables for the secondary earner in a couple . . . 193

A7 RD plots for control variables for the secondary earner in a couple (cont.) 194
A8 Average deducted amount for those who take up the deduction 195
A9 Using gross income of the primary earner instead of taxable income shows no bunching around the kink 196
A10 Characteristics of primary earners with gross income relative to the kink 197
A10 Characteristics of primary earners with gross income relative to the kink (cont.) 198
A11 Total effective marginal tax rates (solid black lines) and decomposition for childless singles and lone parents at kink 1 199
A12 Own deducted amount in 2012 and 2013 200

CHAPTER 1

Introduction and main findings

Labor market shocks can have a large and long-lasting impact on people's careers and lives. Consider a plant closing down in the middle of a recession. The subsequent job loss for individual workers can have severe effects, even though these workers lost their job beyond their own fault. While many people are able to find a new job relatively quickly, for some it can lead to prolonged periods of unemployment, permanently lower wages and consumption (Jacobson et al., 1993) and worse health (Rege et al., 2009). Some research even finds that following a job loss, workers experience an increase in mortality (Sullivan and von Wachter, 2009) and their children perform worse in school (Rege et al., 2011).

Labor market shocks can have many different causes, several of which are explored in this dissertation. An important cause of shocks to workers are structural changes in the economy, such as globalization (Autor et al., 2014) or technological change. Going as far back as the industrial revolution, people have worried about how technology impacts the labor market and replaces some workers. Recently, the development of robots and artifical intelligence has sparked renewed interest in this question (Autor, 2015; Autor and Salomons, 2018; Acemoglu and Restrepo, 2018d).

Chapter 2 examines how recent automation technology, such as robots and AI, impacts individual workers.[1] In this chapter we provide the first micro-level empirical evidence of the effect of automation on a range of worker-level outcomes, including the probability to leave the automating firm, wage income, benefit receipt and self-employment. We observe how much firms spend on automation and identify

[1]This chapter is joint work with James Bessen, Maarten Goos and Anna Salomons. It is based on Bessen et al. (2019).

large increases in these costs as automation events. We then exploit the differences in timing in these events between firms in a differences-in-differences design. We find that for incumbent workers (defined as those with a firm tenure of at least three years) automation at the firm increases the probability to separate from the firm. Firm separation is followed by an increase in time spent in unemployment. Due to the increased incidence of unemployment, workers experience on average a decline in cumulative wage income of around 11% of yearly earnings after five years. We do not find evidence of wage scarring. We find that these earnings losses are pervasive across firm and worker types and only partially offset by benefit systems. Finally, we compare automation events to computerization events, and find no such losses when firms invest heavily in computers.

We contribute to the existing literature with our direct measure of automation at the firm level, which allows us to study the worker impacts of automation where they originate. Second, we develop and implement a methodology exploiting the timing of firm-level automation events for identifying causal effects. Third, we consider automation events across all private non-financial sectors, whereas the existing literature generally only considers a specific automation technology. Fourth, we examine a wide array of worker level outcomes. Finally, we directly compare the current worker-level impacts of automation to those of computerization.

Another important cause of labor market shocks to individual workers is the business cycle, such as when a plant closes down in a recession (Jacobson et al., 1993). Similarly, young workers who enter the labor market in a recession are generally worse off than those who enter in good times (Oreopoulos et al., 2012). Chapter 3 examines the consequences of this shock for Dutch high educated graduates who enter the labor market in a recession between 1996 and 2012.[2] The chapter contributes to the existing literature by examining the effects of graduating in a recession separately for academic and vocational graduates and to consider in detail how job mobility contributes to catching up by looking at how young workers climb the job ladder.

I find that academic graduates suffer strong initial wage effects of 10% for each percentage point decline (around half of a standard deviation) in field-specific employment at graduation. The wage losses gradually decline until they fade out after about five years on the labor market. The initial wage losses for vocational graduates are significantly smaller at close to 6% for each percentage point decline in field-specific employment at graduation. They remain significantly smaller than for university graduates in the first four years. However, wage losses for vocational

[2]This chapter is single authored. It is published as Van den Berge (2018).

graduates remain persistent at about 1% up to at least 8 years after graduation. Employment probabilities for both academic and vocational graduates are negatively affected in the first three to four years on the labor market. While self-employment is not affected for vocational graduates, for academic graduates I find evidence of graduates temporarily substituting regular employment for self-employment in the first years after graduation.

I show that job mobility plays a critical role in recovering from initial wage losses for both academic and vocational graduates who start in a recession. Both groups are more likely to switch firms and sectors, and when they do switch, they gain more than their counterparts who started in a boom. Graduates are more likely to start in firms that pay lower wages in a recession and gradually move to higher paying firms. Both are also more likely to be mismatched in their early career. Interestingly, while switching sectors solves the initial mismatch for academic graduates, vocational graduates remain in sectors that are not typical for their field of study. This could explain the persistent wage losses for vocational graduates.

Institutions and policy can also cause shocks to people's labor market position. Chapter 4 considers how parents adjust their behavior on the labor market to their youngest child going to primary school.[3] Primary school, in addition to teaching children, also functions as both free and compulsory childcare. This is different from most other childcare arrangements studied in the literature, which are often inexpensive or even free, but are not compulsory. We build a theoretical model that shows that the youngest child going to school might have two effects on parental working hours. First, parents who used to take care of their children during school hours experience an increase in free time available and are hence expected to increase their working hours. Second, parents whose children attended paid childcare before going to school might decrease their working hours when their youngest child starts school, because they save on childcare expenses.

Empirically we find significant differences in the responses between men and women. Dutch mothers on average experience an increase in available time of thirteen hours a week when their youngest child goes to school, yet the average number of hours worked per week increases by 0.5 hours after two years. This is an increase of around 3% relative to their mean hours worked. Dutch fathers, who usually already worked full-time, also show a small increase in hours worked of about 0.3 hours, or 0.8% relative to the mean.

[3]This chapter is joint work with Lisette Swart and Karen van der Wiel. It is based on Swart et al. (2019).

We contribute to the existing literature on compulsory schooling and parental labor supply by examining effects both for mothers and fathers. Furthermore, our unique dataset allows us to precisely estimate the effects of compulsory schooling on labor supply. We observe recent cohorts of parents for each month surrounding their youngest child going to school. Finally, the Dutch institutional setting allows us to clearly disentangle the effect of school-going from seasonal effects.

Policy makers are often reluctant or unable to directly intervene in the market processes and choices that can lead to labor market shocks, even if they have negative consequences for some workers. For example, limiting technological progress might help some workers keep their job, but would hamper economic growth. Instead, the policy response often consists in compensating the workers hurt by these shocks, such as through unemployment benefits. This is of course only a temporary answer. Workers who lose their job due to new technology often require new skills to be able find new work. This could be addressed by investing in training. However, it is unclear what the right policy instrument is to promote training. The literature shows that a direct financial instrument, such as a schooling voucher, increases training, but at the cost of a substantial dead weight loss (Schwerdt et al., 2012; Hidalgo et al., 2014). Chapter 5 examines whether a tax subsidy available to all workers instead provides a good incentive for people to invest in training.[4]

Workers in the Netherlands are allowed to deduct training expenses for lifelong learning at their marginal tax rate. To estimate the effect of this deduction on training, we exploit two jumps in the marginal tax rate. These jumps create exogenous variation in the effective costs of lifelong learning for people with very similar income levels. For singles we find heterogeneous effects. For low-income singles we find no effect of the lower costs of lifelong learning due to the jump in the marginal tax rate. However, for high-income singles we find a 10% increase in the probability to file lifelong learning expenditures. We find that these effects are primarily driven by higher-educated middle-aged males. For couples we find small effects for primary earners and no effects for secondary earners.

Chapter 5 builds on the analysis in an earlier paper by Leuven and Oosterbeek (2012), but makes substantial improvements. First, we use more detailed and higher quality data, which allows us to more precisely estimate the effects. Second, we take into account different effects for singles and couples, which turns out to be important for the results. Third, we use a regression-kink design for singles, which is more

[4] This chapter is joint work with Egbert Jongen and Karen van der Wiel. It is based on Van den Berge et al. (2017).

appropriate given the data than the regression-discontinuity design in Leuven and Oosterbeek (2012). Finally, for couples we observe the amount deducted by both partners before and after shifting the deductibles. We show that ignoring this shifting behavior can lead to large spurious estimates for both primary and secondary earners. We also contribute to the literature on financial incentives for lifelong learning, which typically examines direct subsidies rather than tax incentives.

In my dissertation I make extensive use of two important methodological developments. First, the increasing availability of large, high quality administrative data sets covering the entire population. I sometimes combine them with survey data to answer questions that cannot be answered with just administrative or survey data alone. Second, the development of tools to answer causal questions using observational data (Angrist and Pischke, 2009). In this dissertation differences-in-differences and regression discontinuity are applied. Differences-in-differences compares a treatment and control group over time. It relies on the assumption that, while there could be differences between the two groups, the differences do not change over time: both groups would follow a similar trend in the absence of treatment. In this thesis matching of the treatment and control group on observed characteristics is generally applied to ensure that the two groups are as comparable as possible before the treatment. This method is applied in chapters 3 and 4. Regression discontinuity, on the other hand relies on a sharp cutoff in a running variable, such as income or age, that determines whether people are treated or not. This allows a comparison of people just below this cut-off, who are not treated, with people just above the cut-off, who are treated. The assumption for a causal interpretation of this comparison is that all other characteristics, including unobserved characteristics, do not change discontinouosly at the cutoff. This method is applied in chapter 5.

In sum, this dissertation explores how workers adjust to three different labor market shocks. Chapter 2 examines how automation affects individual workers. Chapter 3 explores how young workers adjust to entering the labor market in a recession. Chapter 4 considers how parents adjust their working hours to their youngest child going to school. Finally, chapter 5 examines the effectiveness of one policy which might help workers adjust to shocks: a tax subsidiy for training investments aimed at stimulating workers to learn new skills. Chapter 6 concludes with a short summary of the main findings.

CHAPTER 2

Automatic Reaction: What Happens to Workers at Firms that Automate?*

2.1 Introduction

Advancing technologies are increasingly able to fully or partially automate job tasks. These technologies range from robotics to machine learning and other forms of artificial intelligence, and are being adopted across many sectors of the economy. Applications range from selecting job applicants for interviewing, picking orders in a warehouse, interpreting X-rays to diagnose disease, and automated customer service. These developments have raised concern that workers are being displaced by advancing automation technology. Indeed, opinion surveys from the US and Europe highlight that a majority of individuals are worried about the future of work and expect worsening employment prospects, even as they foresee a positive impact on the economy and on society more generally (Eurobarometer 2017; Pew 2017).

This potential for automation to displace workers is studied in recent labor market models where technology changes the comparative advantage of workers across job tasks (Autor et al. 2003; Acemoglu and Autor 2011; Acemoglu and Restrepo 2018a,d,b; Benzell et al. 2016; Susskind 2017). In these theories, worker displacement at the micro level plays a central role, as machines take over tasks previously performed

*This chapter is joint work with James Bessen, Maarten Goos and Anna Salomons. It is based on Bessen et al. (2019).

by humans. Under certain conditions, such displacement is a possible outcome of automation even in aggregate.

Empirical work on automation has so far mostly focused on robotics – a prime example of automation technology, albeit one that has penetrated only a limited number of sectors – and on more aggregate outcomes.[1] The macro-economic evidence is mixed: Graetz and Michaels (2018) find that industrial robots have had positive wage effects and no employment effects across a panel of countries and industries, whereas Acemoglu and Restrepo (2018c) find that wages and employment have decreased in US regions most exposed to automation by robots. Applying Acemoglu and Restrepo (2018c)'s empirical design to German regions, Dauth et al. (2017) find evidence of positive wage effects, and no changes in total employment. Further, Koch et al. (2019) show that firms that adopt robots experience net employment growth compared to firms that do not, and Dixon et al. (2019) find that a firm's employment growth rises with their robot stock.

Besides macro-economic and firm-level impacts, it is critical to also study automation's effects on individual workers. After all, the absence of displacement in aggregate need not imply the absence of losses for individual workers directly affected by automation. Any such adjustment costs are also of first-order importance for policymakers aiming to diminish adverse impacts out of distributional concerns.

So far, direct empirical evidence on the worker-level impacts of automation is lacking. Existing studies on worker adjustments have used more aggregate sources of variation and do not always focus on causal effects. In particular, Dauth et al. (2018) correlate regional variation in robot exposure with worker outcomes; Cortés (2016) finds that workers switching out of routine-intense occupations experience faster wage growth relative to those who stay; while Edin et al. (2019) show that workers have worse labor market outcomes when their occupation is experiencing long-term decline. To our knowledge, our paper provides the first estimate of the economic impacts on workers when their firm invests in automation technology.

This study makes several contributions. First, we directly measure automation at the firm level and can therefore analyze the worker impacts of automation where they originate: at the automating firms. We do so by linking an annual firm survey on automation costs to Dutch administrative firm and worker databases, allowing us to consider automation across all private non-financial economic sectors. The data are provided by Statistics Netherlands and cover years 2000-2016: we observe

[1]Other papers have looked at cross-sectional features of automation in manufacturing, including Doms et al. (1997) and Dinlersoz and Wolf (2018).

36,490 firms with at least three years of automation cost data, employing close to 5 million unique workers per year on average. Second, we develop and implement a differences-in-differences methodology leveraging the timing of firm-level automation events for identifying causal effects. Third, we consider automation events as they occur across all private non-financial sectors of the economy rather than considering a specific automation technology in isolation, complementing the literature focused on robotics. Fourth, we measure a rich array of outcomes for individual workers for the years surrounding the automation event: this provides insight in how any adjustment costs come about. These outcomes include annual wage earnings, daily wages, firm separation, days spent in non-employment, self-employment, early retirement, and unemployment insurance and welfare receipts. We also look separately at outcomes for incumbent workers (those employed three or more years at the firm prior to the automation event), and for the firm's more recent hires, and consider how impacts differ across worker characteristics. Finally, we directly compare the current worker-level impacts of automation to those of computerization.

We find that automation at the firm increases the probability of incumbent workers separating from their employers. For incumbent workers (those with at least three years of firm tenure), this firm separation is followed by a decrease in annual days worked, leading to a 5-year cumulative wage income loss of about 11 percent of one year's earnings. On the other hand, wage rates are not much affected: that is, we do not see wage scarring for workers impacted by automation. This is in contrast to displacement from mass lay-offs or firm closures, which have been studied in another literature (see Jacobson et al. 1993; Couch and Placzek 2010; Davis and Von Wachter 2011). However, lost wage earnings from non-employment spells are only partially offset by various benefits systems, and older workers are more likely to enter early retirement. Further, automation's displacement effects are found to be quite pervasive across different incumbent worker types as well as firm sizes and sectors. In contrast, we do not find evidence for such displacement from investments in computer technology. This suggests that for incumbent workers, automation is a more labor-displacing force.

This paper is structured as follows. We first introduce our data source, Dutch matched employer-employee data which we link to a firm survey containing a direct measure of automation expenditures. Section 2.3 contains our empirical approach, outlining a definition of automation events and the resulting differences-in-differences estimation framework. Our main results are reported in section 2.4: subsections consider impacts on workers' wage income and its components; additional adjustment mechanisms; robustness checks; and effect heterogeneity. In section 2.5 we directly compare the worker-level impacts of automation to those of computerization. The final section concludes.

2.2 Data

We use Dutch data provided by Statistics Netherlands. In particular, we link an annual firm survey to administrative firm and worker databases covering the universe of firms and workers in the Netherlands. The firm survey is called "Production statistics" ("*Productiestatistieken*") and includes a direct question on automation costs – it covers all non-financial private firms with more than 50 employees, and samples a subset of smaller non-financial private firms.[2] This survey can be matched to administrative company ("*Algemeen Bedrijfsregister*") and worker records ("GBA" and "BAAN" files).

Our data cover the years 2000-2016, and we retain 36,490 unique firms with at least 3 years of automation cost data – together, these firms employ around 5 million unique workers annually on average. We remove firms where Statistics Netherlands indicate that the data are (partly) imputed.[3] We further remove workers enrolled in full-time studies, and those earning either less than 5,000 euros per year or less than 10 euros per day, as well as workers earning more than half a million euros per year or more than 2,000 euros on average per day. For workers observed in multiple jobs simultaneously, we only retain the one providing the main source of income in each year. We use their total earnings in all jobs as the main measure of wage income.

At the worker level, we observe gross wage income as well as days worked – since we do not observe hours worked, we use daily wages as a measure of wage rates. We further observe workers' gender, age, and nationality.[4] A downside to these data is that we neither observe workers' occupations nor their level of education: the former is unavailable entirely, whereas the latter is only defined for a small and selected subset of workers (with availability skewed toward younger and high-educated workers). We further match worker-level data to administrative records on

[2] Firms are legally obliged to respond to the survey when sampled. However, the sampling design implies our data underrepresent smaller firms: we will examine effect heterogeneity across firm size classes to consider how this sample selection affects our overall findings.

[3] In Appendix 2.7.3 we perform robustness checks from several other sample restrictions, including removing firms with outlier employment changes and those undergoing events such as mergers and acquisitions.

[4] In these data, individuals are classified as "Dutch" if they themselves and both of their parents have been born in the Netherlands.

receipts from unemployment, welfare, disability, and retirement benefits. We can track workers across firms on a daily basis, allowing us to construct indicators for firm separation and days spent in non-employment.

The main advantage of the dataset we construct is the availability of a direct measure of automation at the firm level. In particular, "Automation costs" is an official bookkeeping term defined as costs of third-party automation services.[5] While the disadvantage of this measure is that we do not know the exact automation technology being used by the firm, it does capture all automation technologies rather than focusing on a single one, and we measure it at the level of the firm rather than the industry, and across all private non-financial sectors. From discussions with company representatives and automation services providers, we know that these expenditures are related to automation technologies such as self-service check-outs, warehouse and storage systems, data-driven decision making, or automated customer service. Another example are robotics integrator services highlighted (and used as an instrument for robotic technology adoption) in Acemoglu and Restrepo (2018d).

Table 2.1 shows summary statistics on annual automation costs for firms, both in levels, per worker, and as a percentage of total costs (excluding automation costs). This highlights several things. First, almost one-third of firm-year observations has zero automation expenditures. Second, the average automation cost share is 0.44 percent, corresponding to an outlay of around 200K euros annually, or 953 euros per worker. Third, this distribution is highly right-skewed as the median is only 0.15 percent – this skewness persists even when removing observations with zero automation costs.

Table 2.2 further shows how these automation costs and cost shares differ by broad (one-digit) sector. Our comprehensive measure of automation technologies indicates that all sectors have automation expenditures, though there is substantial variation at the firm level both between and within each of these sectors. Average expenditures at the sectoral level range from 220 to 1,636 euros per worker. The highest mean automation expenditures per worker are observed in Professional, scientific, and technical activities, followed by Information and communication, Wholesale and retail, and Manufacturing. Conversely, Accommodation and food serving has the lowest expenditure per worker, followed by Construction, Administrative and support activities, and Transportation and storage. However, there is much variation between firms in the same sector, as shown by the standard deviations of the automation cost

[5] This also includes non-activated purchases of custom software and costs of new software releases, but excludes prepackaged software licensing costs.

Table 2.1: Automation cost share distribution

	All observations			**Automation costs >0**		
	Cost level	Cost per worker	Cost share (%)	Cost level	Cost per worker	Cost share (%)
p5	0	0	0	2,026	54	0.04
p10	0	0	0	3,652	92	0.06
p25	0	0	0	9,537	234	0.14
p50	10,508	257	0.15	27,390	587	0.32
p75	48,000	899	0.47	85,597	1,322	0.68
p90	175,035	2,058	1.05	278,213	2,697	1.37
p95	412,945	3,305	1.69	650,966	4,200	2.13
mean	*192,391*	*953*	*0.44*	*280,713*	*1,391*	*0.64*
N firms × years		240,320			164,707	
N with 0 costs		31%			0%	

Notes: Automation cost level and per worker are reported in 2010 euros, automation cost share is calculated as a percentage of total costs, excluding automation costs. The number of observations is the number of firms times the number of years.

share in total (other) costs. While we do not use either this sectoral or between-firm variation in our empirical identification strategy, we will consider effect heterogeneity across sectors since the nature of automation technologies may be sector-specific.

Table 2.3 reports the same statistics but separately by firm size class, grouped into 6 classes used by Statistics Netherlands: the smallest firms have up to 19 employees whereas the largest have more than 500. Unsurprisingly, automation cost levels rise with firm size: firms with fewer than 20 employees spend around 11K euros annually on automation services, whereas the largest firms spend close to 2.9 million. Less obviously, this table also reveals that automation cost shares increase with firm size, particularly at the very top. The smallest firms have average automation cost shares of 0.40 percent[6], whereas firms with between 20 to 200 employees have a cost share of around 0.44 percent. This increases to 0.51 percent for firms between 200 and 500 workers, and 0.76 percent for firms with more than 500 workers. There is substantial variation within size classes, also.

Figures 2.1 and 2.2 further show how the distribution of automation cost shares and expenditures per worker change over time. Mean automation cost share and outlays per worker are rising in the Netherlands over 2000-2016, from 0.28 to 0.57

[6]The relatively high expenditure per worker for the smallest firm size class is driven by a small number of one-person firms with high automation expenditures – when we eliminate the top 1 percent of observations in terms of automation cost per worker, outlays per worker are monotonically rising in firm size as reported in Table 2.14 in the Appendix.

Table 2.2: Automation costs by sector

Sector	Mean cost level (€) Total	Per worker	Cost share (%) *Mean*	*SD*	Nr of obs *Firms*	*Firms × yrs*
Manufacturing	391,214	986	0.36	0.58	5,655	44,636
Construction	71,150	414	0.2	0.36	4,688	28,757
Wholesale & retail trade	106,259	1,075	0.31	0.80	11,041	75,421
Transportation & storage	257,057	834	0.42	1.07	3,122	21,235
Accommodation & food serving	49,475	220	0.29	0.50	1,292	6,761
Information & communication	409,511	1,636	0.85	2.92	2,655	16,854
Prof'l, scientific, & technical activities	136,437	1,174	1.02	1.76	4,074	23,692
Administrative & support activities	121,301	761	0.49	1.18	3,963	22,964

Notes: Automation cost level in 2010 euros, automation cost shares as a percentage of total costs, excluding automation costs. Total N firms is 36,490; Total N firms × years is 240,320.

Table 2.3: Automation costs by firm size class

Firm size class	Total cost Mean	Cost per worker Mean	SD	Cost share (%) Mean	SD	Nr of obs *Firms*	*Firms × yrs*
1-19 employees	11,135	836	13,255	0.40	1.29	9,850	48,758
20-49 employees	25,287	815	4,152	0.42	1.34	13,777	87,188
50-99 employees	56,336	873	3,975	0.42	0.96	6,291	47,209
100-199 employees	132,573	1,038	5,318	0.44	0.94	3,471	28,748
200-499 employees	372,095	1,440	19,498	0.51	1.11	1,969	17,897
≥500 employees	2,885,712	1,937	13,082	0.76	1.60	1,132	10,520

Notes: Automation cost level in 2010 euros, automation cost shares as a percentage of total costs, excluding automation costs. Total N firms is 36,490; Total N firms × years is 240,320.

percent relative to total other costs, and from 744 to 1,103 euros per worker. All else being equal, this implies that workers' exposure to automation is also rising. Furthermore, besides an increase in the average, there is a fanning out of the distribution with automation cost shares rising faster for higher percentiles.

Lastly, we find that automation expenditures are somewhat correlated with computer investments: these are available from a different, and partially overlapping, firm-level survey. In section 2.5, below, we consider the robustness of our results to excluding firms that have investment events in both technology types within the estimation window, as well as study how the worker impact of computer investment events differs from that of automation.

2.3 Empirical approach

2.3.1 Defining automation cost spikes

The main challenge for empirically identifying the worker-level impacts of automation lies in finding a group of workers who can be used as a control group. A further

Figure 2.1: Firm-level automation cost shares over time

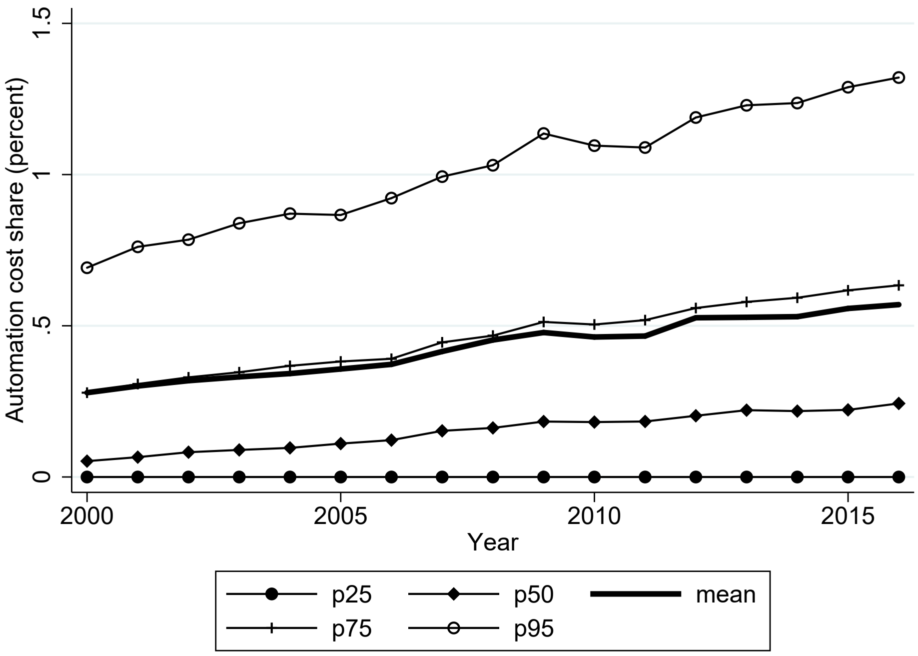

Figure 2.2: Firm-level automation cost per worker over time

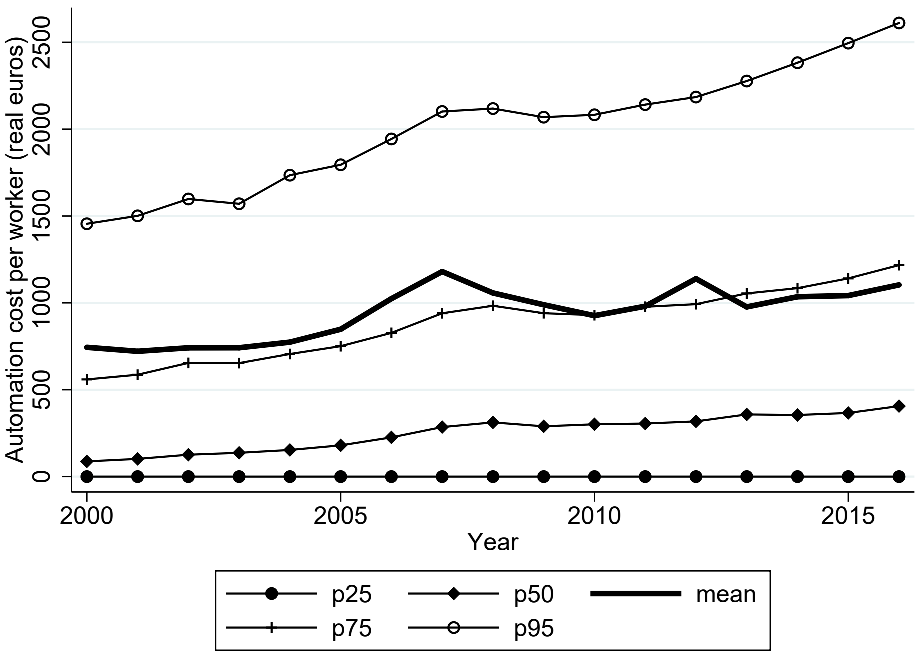

challenge is to distinguish automation events at the firm level, especially when using survey data. Our novel approach for both of these challenges is to use what we term *automation spikes.* In particular, we assume that spikes in automation cost shares at the firm level signal changes in work processes related to automation.

We define automation cost spikes as follows. Firm j has an automation cost spike in year τ if its real automation costs $AC_{j\tau}$ relative to real total operating costs (excluding automation costs) averaged across all years t, $\overline{TC}_j$, are at least thrice the average firm-level cost share excluding year τ:

$$spike_{j\tau} = \mathbb{1}\left\{\frac{AC_{j,t=\tau}}{\overline{TC}_j} \geq 3 \times \frac{\overline{AC}_{j,t\neq\tau}}{\overline{TC}_j}\right\}, \tag{2.1}$$

where $\mathbb{1}\{\ldots\}$ denotes the indicator function. As such, a firm that has automation costs around one percent of all other operating costs for year $t \neq \tau$ will be classified as having an automation spike in $t = \tau$ if its automation costs in τ exceed three percent of average operating costs over years t.

Note that this is a firm-specific measure, intended to identify automation events that are large for the firm, independent of that firm's initial automation expenditure level. As such, this indicator does not mechanically correlate with firm characteristics such as firm size, sector, or capital-intensity. Although we could possibly exploit the size of the automation spike, this is not our specification for a number of reasons. First, there may be measurement error in the survey variable making it more difficult to measure the exact size of a spike. Second, we use the automation costs survey variable to flag automation events, but other (indirect) costs may be incurred which are not directly surveyed: as such, our baseline approach identifies automation events without taking a strong stance on their exact size. In Appendix 2.7.3.5, we report several robustness checks, including changing the automation spike definition and varying the spike threshold.

The existence of these automation cost spikes would be consistent with a literature on lumpy investment (Haltiwanger et al. 1999; Doms and Dunne 1998). In fact, such spikes occur when the investment is irreversible and there are important indivisibilities. Under uncertainty, irreversibility creates an option value to waiting (Pindyck 1991; Nilsen and Schiantarelli 2003); whereas indivisibilities can arise from fixed adjustment costs (Rothschild 1971) – together, this implies investment occurs in relatively infrequent episodes of disproportionately large quantities. It is plausible that investments in automation meet these two criteria: major automation investments likely both include substantial irreversible investments (for example in terms of worker training or from developing custom software) as well as involve fixed adjustment costs from reorganizing production processes.

2.3.2 Summary statistics on automation cost spikes

We now document the existence and frequency of automation spikes by firm and sector. In order to identify spikes, we need at least three years of automation cost data at the firm level: this is the sample of 36,490 firms described above.

Table 2.4 shows that around 70 percent of firms never spike, whereas the remaining 30 percent spike at least once over the 17 years of observation. Note that non-spiking firms do not necessarily have zero automation costs: it is just that their automation expenditures do not fluctuate much as a percentage of total costs, implying they do not undergo large automation events as we define them. Out of the firms that do have such an event, the large majority spikes only once over 2000-2016, although some spike twice and up to five times at most. Automation spikes are observed across all sectors, as Table 2.5 highlights. However, a higher share of firms in Information and communication experience such an event compared to firms in Construction or in Accommodation and food serving.

Figure 2.3 shows what automation spikes look like on average across firms where spikes are observed. This is constructed by redefining time t as the number of years relative to the spike in period τ, i.e. $t \equiv \mathit{year} - \tau$, such that all spikes line up in $t = 0$. When firms spike multiple times, we only include the largest spike. Figure 2.3 is not using a balanced panel of firms: rather, all 10,476 spiking firms are observed in $t = 0$, and the number of observations for other years depends on when the spike took place[7], and on how often the firm enters in the automation survey. Nevertheless, we see a clear spike pattern.

Figure 2.4 restricts the sample of firms with spikes in $t = 0$ to those firms that are observed in all years $t \in [-3, 4]$, as these are the treatment group firms we will actually use in the empirical design explained below.[8] Figure 2.4 shows that automation events are quite cleanly identified: these events are not preceded by a substantial lead-up of automation spending relative to total costs, nor is there evidence of much slow tapering off afterwards. Rather, automation spike years stand out as years when the firm made a large (relative to its normal automation expenditure share) investment in automation.

Figure 2.5 repeats Figure 2.4 but for the implied level of automation expenditure per worker, showing that the average firm-level automation spike amounts to an investment of close to 1,900 euros per worker, compared to a usual level of around

[7]For example, if the spike occurred in the first calendar year of data, there are no observations for $t < 0$; if it took place in the last calendar year, there are no observations for $t > 0$.

[8]See Appendix 2.7.1 for details on sample construction.

440 euros in years close to the spike. Figures 2.4 and 2.5 are both weighted by firms' employment size: as such, they reflect the exposure to automation for the average treated worker in our sample.[9]

2.3.3 How do automating firms differ?

A potential control group for workers in automating firms are workers in firms that are not automating. However, here we show that these groups are not comparable. Table 2.6 first considers how the average automation expenditures compare across these two groups. This reveals that firms with automation events have higher average levels of automation expenditures, whether expressed in absolute terms, or relative to the number of workers, or as a share in total costs. These differences are considerable: firms with automation events spend around twice as much on automation per worker or relative to total operating costs.[10]

Importantly, firms that make large automation investments have faster employment growth compared to firms that do not have automation spikes. This is shown in Figures 2.6 and 2.7 which respectively plot firm-level log employment and wagebill trajectories, for a balanced sample of firms existing over the entire 17-year period. These stark descriptive differences in trajectories between automating and non-automating firms are consistent with findings in Koch et al. (2019), and in part motivate our empirical design, outlined in the next section.

2.3.4 Empirical design

We now outline our empirical design to leverage the observed automation cost spikes for identification. Our specification only considers incumbent workers who are employed in firms that spike at some point over 2000-2016. We define incumbent workers as workers with at least 3 years of firm tenure. This by and large captures workers with permanent contracts and hence workers who have a stable working relation with the firm.[11] This is important because identification requires that

[9]In Figures 2.21 and 2.22 in the Appendix, we show that the same patterns hold when considering an entirely balanced sample of treatment firms where we observe automation cost share information in every single year.

[10]Further, Table 2.15 in the Appendix shows how spiking firms' time-invariant characteristics differ from non-spiking ones: the main finding is that firms that experience automation spikes are larger.

[11]Dutch labor law during almost our entire data period ensures that temporary contracts are of a maximum duration of 3 years, implying that workers with 3 years of tenure are very likely to have permanent contracts. On average across firms in our data, 64 percent of workers are incumbents (where the median is 70 percent).

Table 2.4: Firm-level automation spike frequency

Spike frequency	**N firms**	**% of N firms**
0	26,014	71.3
1	8,411	23.1
2	1,764	4.8
3	267	0.7
4	30	0.1
5	4	0.0
Total	36,490	100.0

Notes: Spike frequency is defined as the total number of spikes occurring over 2000-2016. The total number of firms is 36,490.

Table 2.5: Automation spike frequency by sector

Sector	**N firms**	**N firms with spike**	**Spike frequency (%)**
Manufacturing	5,655	1,606	28.4
Construction	4,688	1,143	24.4
Wholesale & retail trade	11,041	3,004	27.2
Transportation & storage	3,122	937	30.0
Accommodation & food serving	1,292	329	25.5
Information & communication	2,655	1,023	38.5
Prof'l, scientific, & technical activities	4,074	1,293	31.7
Administrative & support activities	3,963	1,141	28.8

Notes: A spiking firm has at least once automation spike over 2000-2016. The total number of firms is 36,490, the total number of spiking firms is 10,476. Spike frequency is the ratio of spiking firms over total firms by sector.

Table 2.6: Automation expenditures by firm type

	Mean automation cost:		
Firm type	level	per worker	share (%)
No automation spike	245,066	1,389	0.62
$\geq$1 Automation spike	359,797	2,547	1.29

Notes: Total N firms is 36,490.

workers are not self-selected into the firm in anticipation of an automation event occurring in the near future.[12] This reasoning is similar to the focus on incumbent workers in the mass lay-off literature (e.g. see Jacobson et al. 1993; Couch and Placzek 2010; Davis and Von Wachter 2011).

[12]In section 2.4.4, we also estimate impacts for the group of workers with less than three years of firm tenure prior to the automation event. Causal identification of the treatment effect for this group is more difficult as they may have been hired in anticipation of the automation event. We therefore analyze them separately, and generally put more stock in our results for incumbent workers.

Figure 2.3: Automation cost share spikes

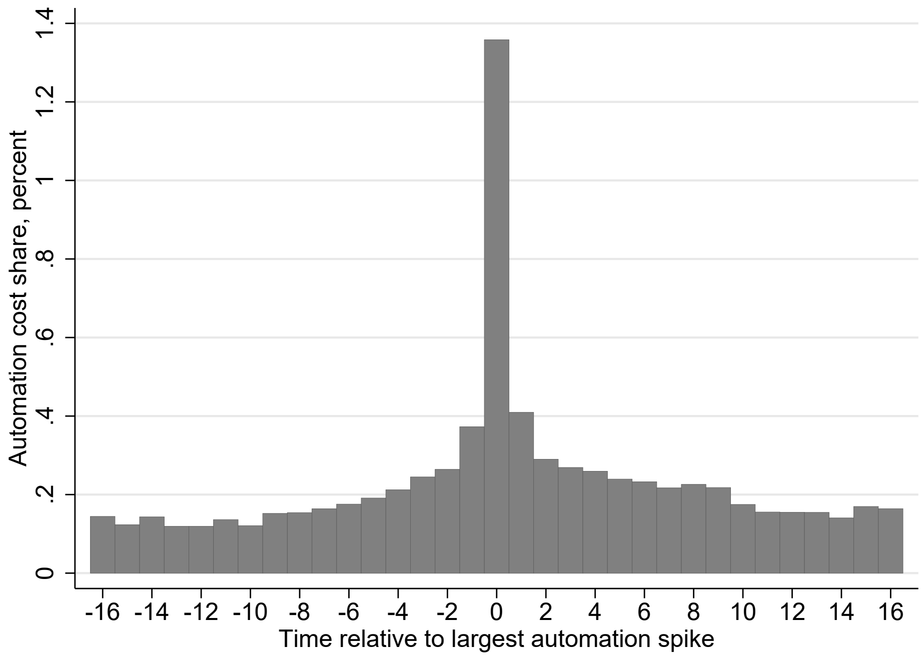

Notes: Unbalanced panel of firms, N=10,476 in $t = 0$.

Figure 2.4: Automation cost share spikes for treated firms

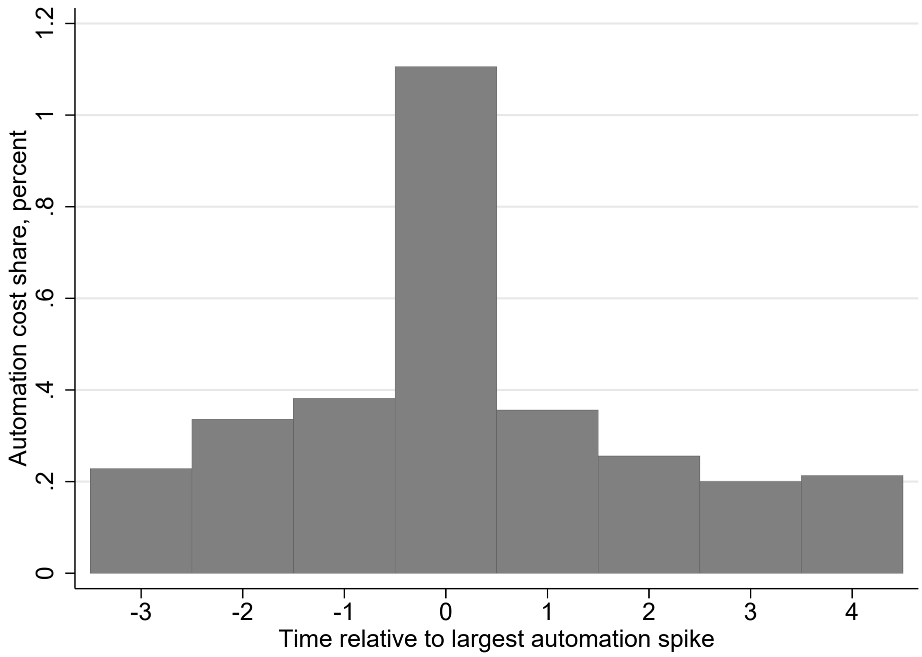

Notes: N=2,446 in $t = 0$.

Figure 2.5: Automation cost level per worker for treated firms

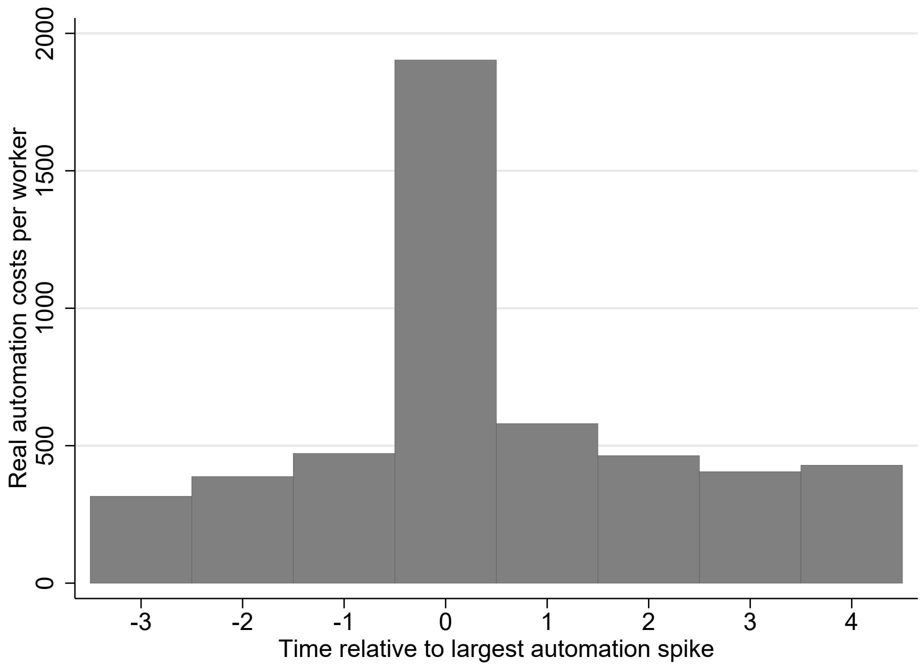

Notes: N=2,446 in $t = 0$.

We define the group of treated workers as those with 3 or more years of firm tenure at $t - 1$ in treatment group firms, i.e. firms that spike in $t = 0$ and are observed in all years $t \in [-3, 4]$. Treated workers are further divided into cohorts by the calendar year in which their firm spikes. Specifically, given that our sample covers calendar years 2000 to 2016, the earliest cohort of treated workers are those employed between 2000 and 2002 at a firm that spikes in 2003. Similarly, the last cohort of treated workers are those employed between 2008 and 2010 in firms that spike in 2011.[13]

For each cohort of treated workers, we then define a control group of workers with at least 3 years of firm tenure at $t - 1$ and who are, at $t - 1$, employed in firms that spike in $t + 5$ or later.[14] For example, the control group for the earliest cohort of treated workers are workers employed between 2000 and 2002 at a firm that spikes in 2008 or later. Similarly, the control group workers for the last cohort of treated workers are those employed between 2008 and 2010 at the same firm that spikes in 2016. Finally, we exclude both treatment and control group firms with multiple

[13] See Appendix 2.7.1 for more details on sample construction.

[14] We only require control group workers to be at a firm j that spikes at $t + 5$ or later to stay at firm j from $t = -3$ until $t = -1$. Hence, they do not have to be employed at firm j when firm j actually spikes in year $t + 5$ or later.

Figure 2.6: Log employment for firms with and without automation events

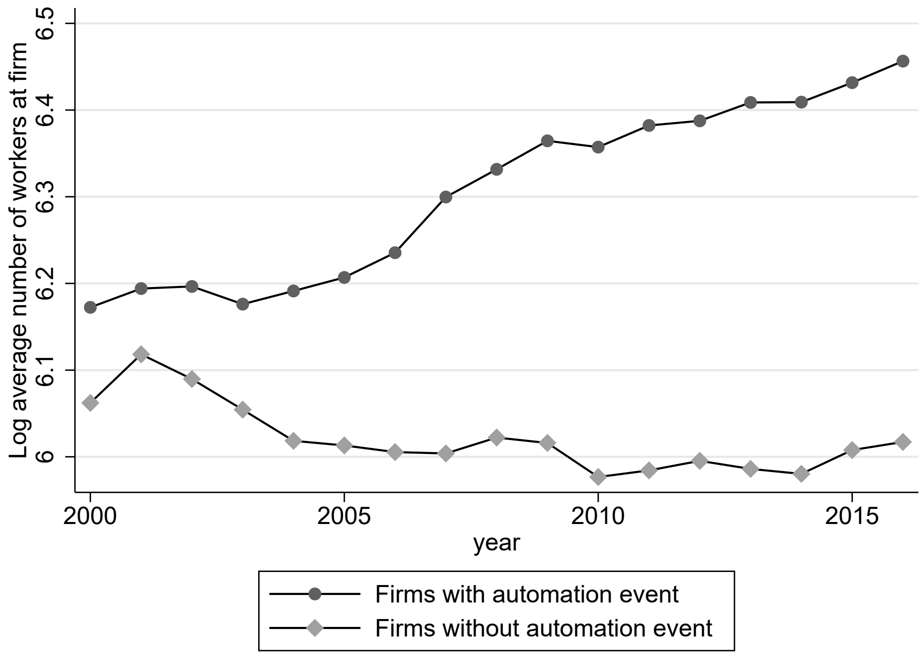

Notes: Balanced sample of 399 firms with and 623 firms without an automation event.

spikes in the estimation window such that estimates of pre-trends and treatment lags are not contaminated, but our results are similar when not imposing this restriction.

By defining treatment and control group workers from firms that spike at least once (i.e. excluding workers from firms that never spike in our control group), our specification strictly exploits differences in event timing rather than also using event incidence for identification. As such, we assume that from the perspective of incumbent workers, the timing of automation cost spikes is essentially random conditional on observables.[15] Another way to think about our approach is that we match workers on the firm-level outcome of making large investments in automation technology at some point in time. Only exploiting spike timing (rather than also spike incidence) across firms is important since in section 2.3.3 we showed that firms with automation events are on very different employment and wagebill trajectories: as such, the employment trajectories for workers employed at firms without these events are not an appropriate counterfactual.

Our use of timing differences across firms is in the spirit of a recent literature exploiting event timing differences in other contexts (see e.g. Duggan et al. 2016;

[15]In Appendix 2.7.2, we use a k-fold cross-validation prediction to show that spike timing is difficult to predict based on observables, increasing our confidence that event timing is plausibly random from the perspective of a firm's incumbent workers.

Figure 2.7: Log wage bill for firms with and without automation events

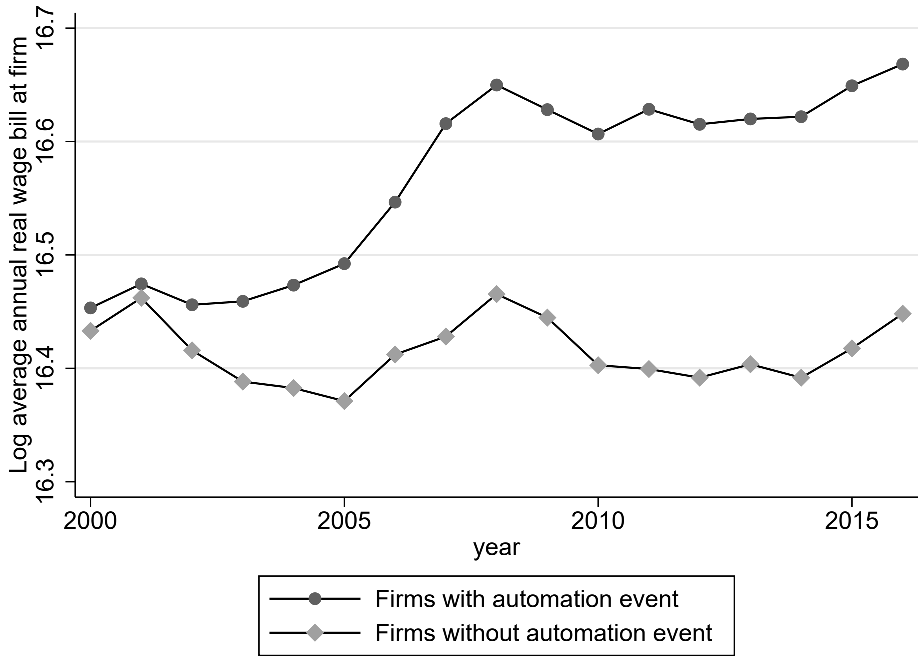

Notes: Balanced sample of 399 firms with and 623 firms without an automation event.

Fadlon and Nielsen 2017; Miller 2017; Lafortune et al. 2018). In the context of automation, our identification relies in part on the nature of major automation events. Indeed, as argued above, because these investments typically involve both uncertainty about the payoff and irreversible investments, they can create substantial option value to waiting to invest. This means that small differences in the payoffs to automating can generate substantial differences in the timing of investment.[16] This sensitivity implies that small, idiosyncratic differences can change the exact timing of automation events across firms. Consequently, workers employed at cohorts of firms that spike a few years apart should be on similar trends, and can thus serve as a counterfactual.

We use a Differences-in-Differences (DiD) specification for each cohort of treatment and control group workers, with the data stacked across cohorts:

$$y_{ijt} = \alpha + \beta treat_i + \sum_{t \neq -1; t=-3}^{4} \gamma_t \times I_t + \sum_{t \neq -1; t=-3}^{4} \delta_t \times I_t \times treat_i + \lambda X_{ijt} + \varepsilon_{ijt}, \quad (2.2)$$

[16] For example, Bessen (1999) finds that a 6 percent payoff difference generated a decade difference in when firms chose to switch from mule-spinning to ring-spinning in the British textile industry.

where i indexes workers, j firms, and $t \in [-3, 4]$ the number of years relative to the timing of the automation spike.[17] y_{ijt} is the outcome variable (such as total wage earnings, annual days in non-employment, wages conditional on working, and firm separation), and $treat_i$ is a treatment indicator, equal to 1 for worker i if their firm is experiencing an automation spike at time $t = 0$, and 0 otherwise. Further, I_t are indicators for time relative to the spike year, with $t = -1$ as the reference category. Lastly, X_{ijt} are controls: these are a set of worker characteristics (age and age squared, gender, and nationality); sector and size class of the spiking firm; as well as fixed effects for years. In our baseline specification, we replace $\beta treat_i$ with individual fixed effects[18] – this also absorbs non-time varying controls (gender, nationality, firm size and sector). We cluster standard errors at the level of the firm where treatment occurs: that is, all workers employed at the same firm in $t - 1$ are one cluster.

In equation 2.2, the parameters of interest are δ_t: these estimate period t treatment effect for $t \geq 0$, relative to pre-treatment period $t = -1$. As with all DiD models, identification requires parallel trends in the absence of treatment, or that $\delta_t = 0$ for all $t < 0$. Our event timing strategy is intended to support the assumption that worker outcome variables would have followed similar trends in the absence of treatment.

We can further strengthen the assumption of parallel trends by matching on worker and firm observables to ensure that $\delta_t = 0$ for all $t < 0$ (Azoulay et al. 2010). In our baseline specification, we match treated and control group workers on pre-treatment annual real wage income, separately by sector and calendar year. While the match is exact for calendar year and sector, we use coarsened exact matching (CEM, see Iacus et al. 2012a; Blackwell et al. 2009) for pre-treatment income. To this end, we construct separate strata for each 10 percentiles of real wage income, as well as separate bins for the 99th and 99.9th percentiles, in each of the three pre-treatment years $t = -3, -2, -1$. We then match treated workers to control group workers for each of these income bins, while additionally requiring them to be observed in the same calendar year, and work in the same sector one year prior to treatment. We include calendar year and sector matching to ensure we are not capturing sector-specific business cycle effects, or other unobserved time-varying shocks affecting workers based on their original sector of employment. As such, each

[17] Our results are robust to changes in the number of estimated post-treatment periods, which in our setting also changes the set of control group firms.

[18] Except when we estimate individual workers' hazard of leaving the firm.

treated worker is matched to a set of controls from the same calendar and sector and belongs to the same pre-treatment earnings percentile bin. This procedure results in 30,247 strata for incumbent workers[19], and in doing so can match 98 percent of treated incumbents (using 93 percent of control group incumbents).[20]

After matching, our sample contains $1,046,995$ distinct incumbent workers in treatment and control groups. Of those incumbent workers, $102,599$ are treated. Given we observe each of these individuals for 8 years, this results in $8,375,960$ observations. Our estimation sample of firms for identifying these treated and control group workers contains $5,970$ unique firms, all of which experience an automation spike at some point over the period and are observed for at least 8 consecutive years. Workers employed at $2,429$ of these firms are treated, and workers at $4,543$ firms serve as controls at least once.[21]

2.4 The impact of automation on incumbent workers

Here, we consider how incumbent workers are impacted by an automation event at their firm. In the first section, we study impacts on wage income and its components: changes in firm separation and non-employment on the one hand, and changes in daily wages on the other. The next subsection then considers adjustment margins for displaced workers (in particular, sectoral switching, early retirement, and self-employment) and to what extent income impacts are offset by various benefit payments. The third subsection performs a range of important robustness checks on our main results, and the final subsection considers effect heterogeneity in these impacts.

[19]Note that some strata may not contain any treated workers, in which case they are irrelevant for estimation. Further refinement of these strata does not change our results, although it leads to a smaller percentage of treated workers being matched.

[20]Further support for the parallel trends assumption is given in Table 2.18 in the Appendix. This table compares observables across the treatment and control groups using matching weights. The two groups are closely matched on a wide range of variables for both firms and workers. In Appendix 2.7.3.3, we also show that our results are robust to additionally matching on pre-treatment employment growth at the firm level, incumbent worker firm tenure, and firm size.

[21]Firms can serve as control group firm more than once for different treatment group cohorts. For example, a firm spiking in 2016 might be a control group firm for a firm spiking in 2005 and for another firm spiking in 2008. In addition, some firms can serve as control group firm and as treatment group firm. For example, a firm spiking in 2010 could be a control group firm for a firm spiking in 2005, while also serving as a treatment group firm in 2010.

2.4.1 Impacts on wage income, firm separation, and non-employment

Figure 2.8 shows the impact of automation events on annual real wage income for incumbent workers, that is, workers with at least three years of firm tenure. We scale each individual worker's real wage income by their real wage income level one year before the automation spike, to obtain relative impacts.[22] That is, it shows estimates of equation 2.2, where the outcome variable is annual real wage income over $t = -1$ annual real wage income.

Figure 2.8 shows there are no pre-trends in wage earnings. The estimates highlight that incumbent workers lose income as the result of an automation event. Indeed, in the automation year, incumbent workers lose some 0.9 percent in wage income and this effect increases over time, cumulating to 10.7 percent in total after five years. Estimates are statistically significant for all treatment and post-treatment years. Given that annual earnings grow by 1.6 percent annually on average, this reflects a non-negligible loss compared to usual earnings trajectories. In levels, this corresponds to approximately 323 euros lost for the average worker in the treatment year (0.9 percent of the average pre-treatment income of 35,885 euros)[23]; and 3,839 euros after 5 years in total ($0.107 \times 35{,}885$). This suggests automation leads to displacement for workers: compared to workers employed at firms who automate later, workers employed at currently automating firms experience income losses.

To further scale these costs, we can calculate the incumbent worker income losses resulting per euro of automation expenditure per worker during an automation event. These estimates should be interpreted as an upper bound, since we do not necessarily observe all outlays associated with the automation event we identify. As reported above, during an automation event, automation expenditures are on average 1,912 euros per worker. Since the average worker loses 323 euros in the automation year, and 3,839 euros after 5 years in total, 1 euro invested per worker leads to a loss of around 0.30 euros of per incumbent worker in the automation year, and 2 euros after 5 years.

These annual wage impacts may be driven by changes in days worked following firm separation, changes in daily wages conditional on being employed, or a combination of both. We now consider the first of these adjustment margins:

[22] This is preferable to log income impacts since this would eliminate zeros: this approach is also taken in e.g. Autor et al. (2014). In Figure 2.23 in the Appendix we additionally show impacts in levels.

[23] See Table 2.18 in the Appendix.

Figure 2.8: Annual real wage income, relative to $t = -1$

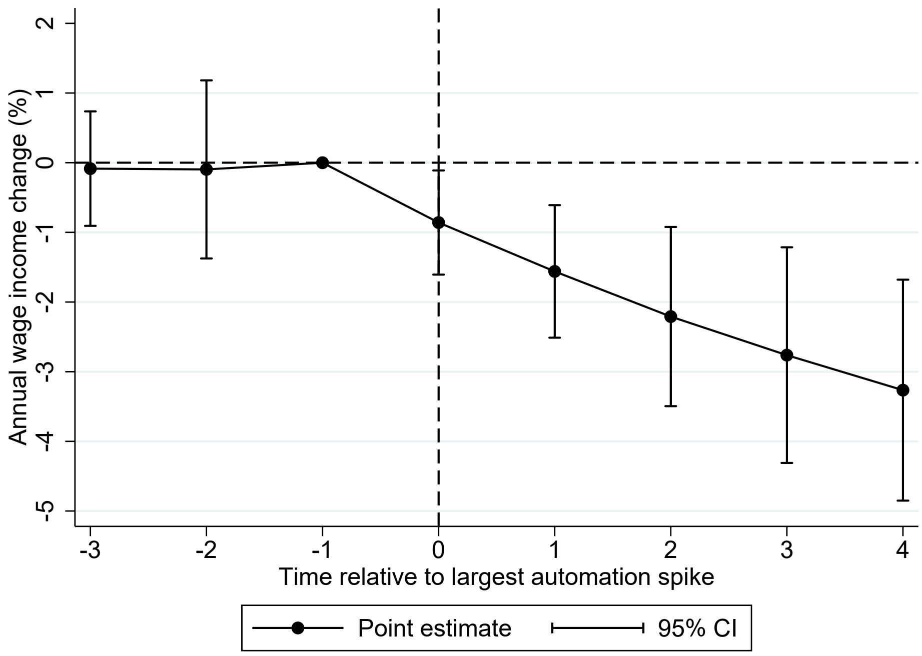

Notes: N=8,375,960. Whiskers represent 95 percent confidence intervals.

do workers separate from their firms as a result of automation, and does this lead to non-employment spells?[24]

Figure 2.9 considers the impact of automation on displacement in its most literal sense: does automation result in incumbent workers leaving the firm? This figure presents estimates from equation 2.2, where the dependent variable is worker-level hazard of separating from their pre-treatment employer.[25] All coefficients have been multiplied by 100, such that the effects are in percentage points. This highlights that workers' probability of separating from their employers following an automation spike is rising over time compared to control group workers.

Specifically, in the automation year, the separation hazard for incumbent workers is 2.1 percentage points higher (though this estimate is statistically insignificant), where the (matched) control group incumbent separation probability is 9.6 percent. After five years, incumbents have a statistically significant 3.6 percentage point

[24]Any adjustment in days worked can in principle come from either the intensive or extensive margin: that is, workers may work fewer days with their current employer, or separate from their employer and experience a non-employment spell before finding re-employment. However, we do not find any evidence of intensive margin changes in non-employment– as such, any change found here reflects adjustments along the extensive margin.

[25]Because the dependent variable is a hazard rate, this model does not include worker fixed effects (unlike estimates for all other dependent variables).

higher firm separation hazard: this is a substantial 41 percent rise compared to the average corresponding hazard among control group workers of 8.8 percent.

It is noteworthy that worker displacement does not occur instantly: rather, displacement effects arise over time. There are various (and non-mutually exclusive) possible explanations for this. For one, these patterns are consistent with incumbent workers having open-ended contracts, making it costly to fire them. Further, these gradual changes could in part also result from a time delay in the effective implementation of automation technologies relative to the cost outlay, or because it takes time for workers and firms to learn about changes to their match quality under the new technology. Gradual displacement is in contrast to the patterns seen during mass lay-off events, where at least 30 percent of the firm's incumbent workforce is laid off at once (see Davis and Von Wachter (2011) for an overview). These gradual changes in firm separation are also what underlie the increasing effects for total wage income seen in Figure 2.8: estimates for any one post-event year reflect the impacts of all worker cohorts having left up until that point.[26]

All in all, this shows incumbents are more likely to separate from their employer as a result of automation. Indeed, the resulting increase in separation is substantially higher relative to what they experience in the absence of an automation event. Although this increase in firm separation implies automation leads to displacement for the firm's incumbent workers, it need not translate to income losses if these workers find re-employment quickly (and at similar wage rates): we now turn to impacts on non-employment.

Results are shown in Figure 2.10, where we define the dependent variable in equation 2.2 as the annual number of days spent in non-employment. Note that incumbents are by definition employed in the three years prior to the automation event – although they need not work full-time and may change their annual days worked, their number of days in non-employment does not evolve much prior to $t = 0$. Starting in the year in which the automation spike takes place, however, their days worked gradually decreases. In particular, non-employment increases by 1.3 days in the automation event year, and this increases to around 5.7 days annually five years after the automation event, with a total cumulative increase in non-employment of 18 days compared to the control group. By comparison, in the event year, matched control group incumbents spend around 5.7 days in non-employment on average,

[26]Put differently, our estimates combine time and cohort effects: since workers are still leaving the automating firm at increased rates several years after the event, and assuming that displaced workers do not all adjust within the span of a year, the average treatment effect consists of cohort effects that cumulate over time.

suggesting automation produces an increase of 22 percent in non-employment in the automation year itself. The cumulative five-year impact also corresponds to a 22 percent increase relative to the five-year cumulative non-employment duration (82 days) experienced by control group incumbents.

Figure 2.9: Firm separation hazard

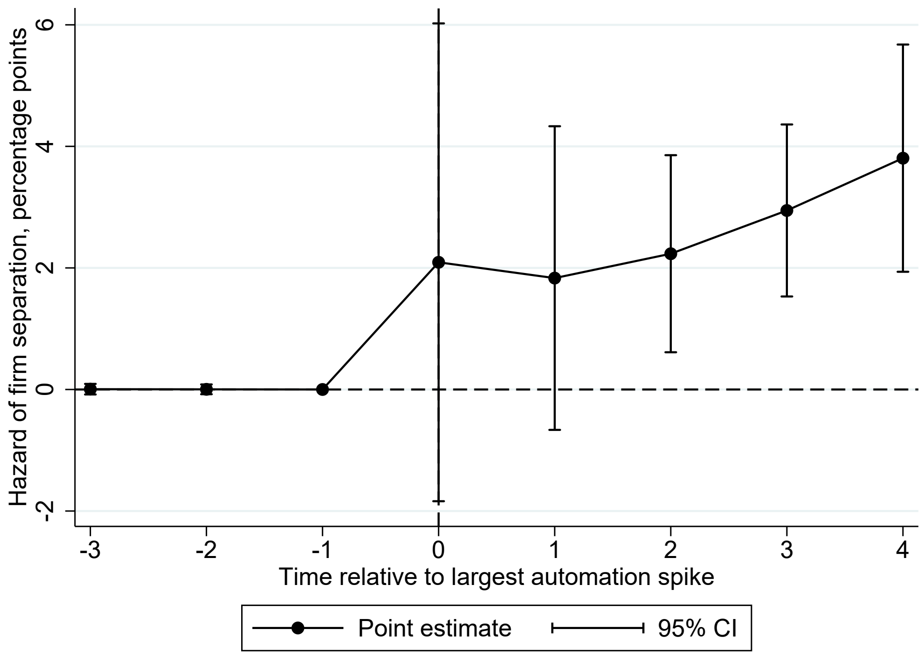

Notes: N=8,375,960. Whiskers represent 95 percent confidence intervals. Effects at $t = -2$ and $t = -3$ are zero by definition since incumbents are at the firm for three years before $t = 0$.

Figure 2.10: Annual number of days in non-employment

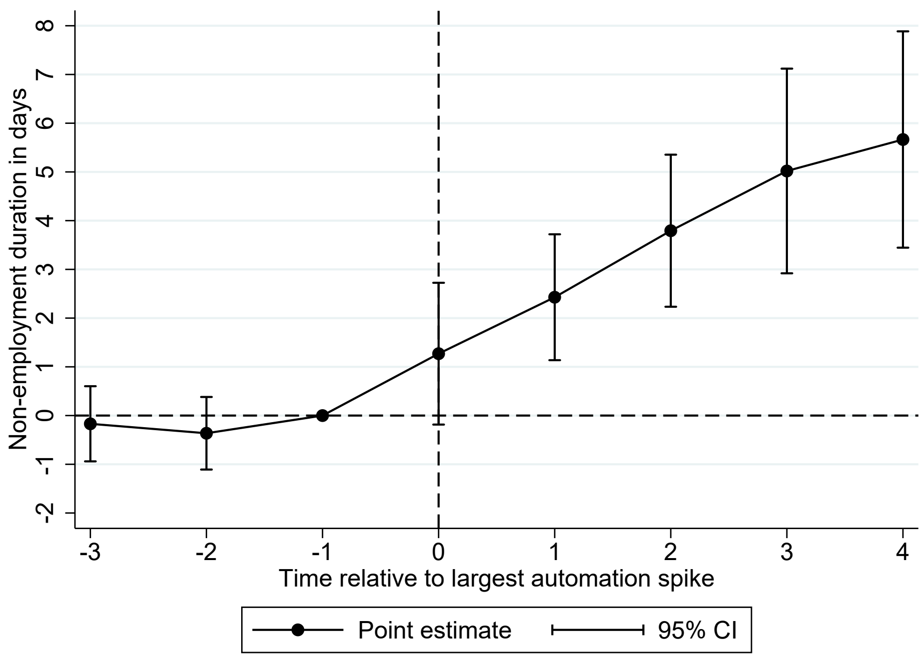

Notes: N=8,375,960. Whiskers represent 95 percent confidence intervals.

We now turn to the daily wage impacts of automation. In Figure 2.11, we consider the effect of automation on log daily wages (conditional on employment) of incumbents. Recall that we do not observe daily hours worked in our data: changes in daily wages can therefore result from changes in hourly wages and/or changes in daily hours worked. We do not find any statistically significant effects, and the point estimates are also economically very small.[27] These findings are in contrast to long-term wage scarring found in mass lay-off studies: also in the Dutch context (and with the same administrative records we are using), such wage scarring has been found (see Deelen et al. 2018; Mooi-Reci and Ganzeboom 2015). This suggests the adjustment costs arising from automation events are more transitory, even for workers with relatively long firm tenure. This need of course not be because of the nature of the automation event, but could in part be because the workers affected by automation are different from those affected by firm closures, in ways that allow them to better deal with job transitions. The absence of daily wage effects also implies that the income losses incumbent workers experience are entirely accounted for by firm separation followed by non-employment spells.

[27] Of course, these wage effects combine effects across job leavers and job stayers, which may cancel out on average – we therefore also estimate our daily wage models separately for these two groups, relative to control group workers of job leavers and job stayers, respectively. Also for these two groups separately, we find only economically small and statistically insignificant impacts.

Figure 2.11: Log daily wages

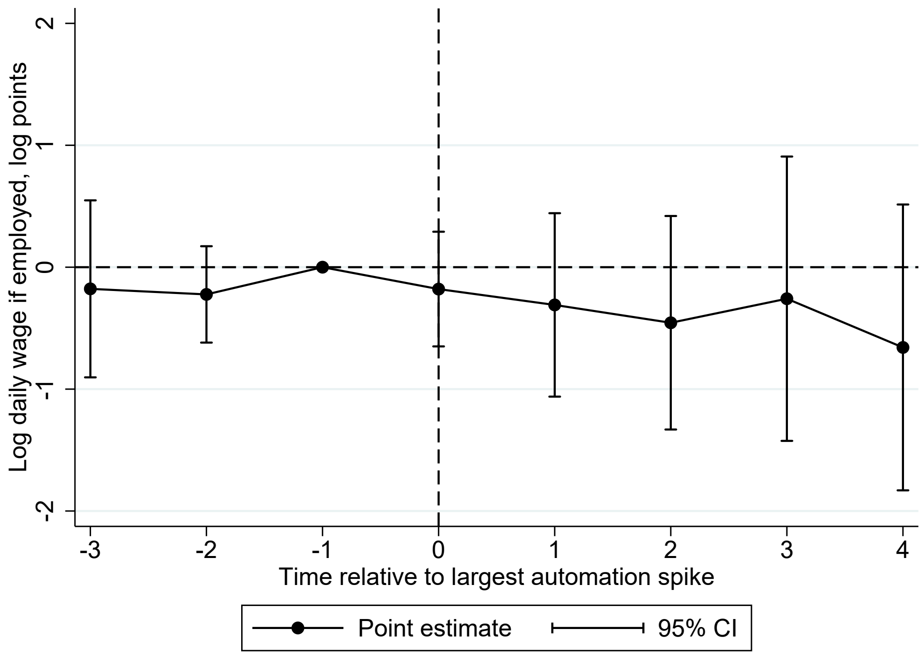

Notes: N=8,094,856. Whiskers represent 95 percent confidence intervals.

Overall, the results in this section show that automation leads to displacement for individual incumbent workers. This of course does not imply automating firms are necessarily displacing workers on net. In Section 2.3.3, we already documented that automating firms expand employment more rapidly than non-automating firms. While our empirical design cannot claim causal identification of automation's employment effects at the firm level (as the automation event and its timing are clearly endogenous from the firm's perspective), in a forthcoming paper we show descriptively that firms do appear to be labor-saving on net around automation events (Bessen et al. 2020). This suggests that labor-saving automation is one way firms may gain a competitive advantage that allows them to expand the size of their operations – including employment– in the longer run. However, these events are accompanied by adjustment costs borne by incumbent workers.

2.4.2 Where do automation-affected workers go?

So far, we have shown that incumbent workers experience income losses as a result of automation in their firm, cumulating to 3,839 euros on average per incumbent worker in total over the five-year post-treatment window. These losses are entirely driven by higher firm-separation probabilities accompanied by non-employment spells. This raises the question where workers affected by automation events go: in this section, we consider a range of outcomes and adjustment mechanisms. First, we will consider to what extent workers impacted by automation are switching to firms in other industries, and to firms of different sizes and different average wages. Next, we study to what extent various benefit systems are compensating for lost wage income. Lastly, we consider whether treated workers are more likely to be observed in early retirement or self-employment compared to the control group.

Figure 2.12 estimates our empirical specification with the probability of switching two-digit industries as the dependent variable. This shows that workers impacted by an automation event are 5 percentage points more likely to switch industries after five years. While the estimates are relatively noisy, this is a 17 percent increase compared to the 30 percent probability of switching industry for control group workers over the same time period.[28] That is, the additional displacement produced by automation is translating to some increased industry switching.

Besides industry switches, we do not find economically sizable or statistically significant changes in terms of workers' average or median firm wage, firm size, or firm automation expenditure.[29] This implies that automation-affected workers are not structurally moving to firms that pay different wages, are of different sizes, or are differently automation-intense.

Figure 2.13 considers the impact of automation on workers' total benefit receipts (in annual real euros), comprised of unemployment benefits, disability benefits, and welfare payments, as well as the separate contributions from these three different sources.

We find that incumbent workers receive additional benefit income following an automation event: after five years, the cumulative amount received is around 514 euros on average, implying that only 13 percent of the negative wage income impact is offset. This finding is comparable to that in other worker displacement events, where typically only a small part of the negative impact is compensated by social security (Hardoy and Schøne 2014).[30]

[28]As shown in the figure results are very similar for one-digit industries, though not statistically significant at the 5 percent level.

[29]For each of these (firm wage, firm size, and firm automation expenditure), we measure the dependent variable in $t = -1$, to keep overall changes in firm characteristics from impacting the estimate. As such, for workers remaining with their pre-event firm, or those switching to a firm that had the same size, wage, or automation expenditure in the pre-treatment year, the change in the dependent variable is zero.

[30]In part, this is by law: unemployment benefits in the Netherlands have a replacement rate of 75 percent in the first two months of unemployment, which then decreases to 70 percent. Further,

Figure 2.12: Probability of switching industries

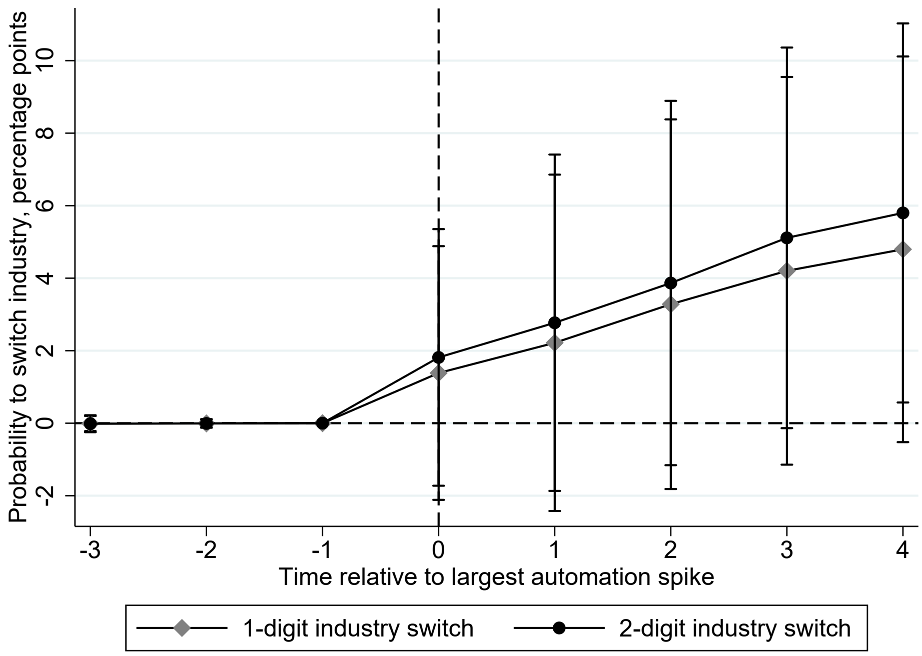

Notes: N=8,375,960. Whiskers represent 95 percent confidence intervals.

Figure 2.13 further shows that all of the benefit payments for incumbents arise from unemployment insurance: this is expected, as unemployment benefit eligibility is very high among workers with at least three years of firm tenure.[31] Consistent with high unemployment benefit eligibility, we do not see any contribution from welfare payments for incumbents. Lastly, disability benefits[32] are actually slightly decreasing over time. Since all incumbents were previously employed, this implies some were working part-time and receiving benefits for a partial disability prior to the automation event. The decline in disability insurance receipts is driven by incumbents who find re-employment – that is, they no longer receive these benefits with their new employer.

Besides benefits and welfare payments, displaced workers could also adjust by entering self-employment: since self-employment income is not observed in our data,

there is a maximum ceiling, such that workers with higher wages earn lower replacement rates than the 70 or 75 percent maximum.

[31]For most of the observation period, eligible workers in the Netherlands are entitled to up to 38 months of unemployment benefits following job loss. In the last year of observation (2016), this has been decreasing: currently, eligibility is 24 months.

[32]Disability benefits in the Netherlands cover impairment whether full or partial, and whether temporary or permanent, and replace up to 70-75 percent of workers' past wages. Benefits are financed by employers without worker contributions, and there is a long history of the use of these schemes in the Netherlands as hidden unemployment (Koning and Lindeboom 2015).

Figure 2.13: Annual real benefit income

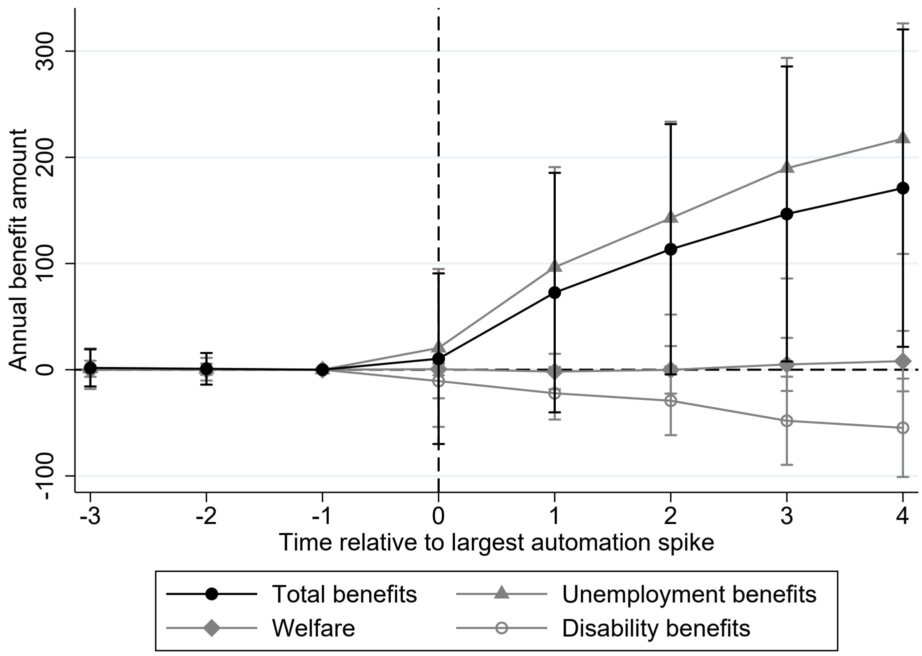

Notes: N=8,375,960. Total benefits are the sum of received unemployment and disability benefits, and welfare. Whiskers represent 95 percent confidence intervals.

we may be overestimating the income losses workers experience. Indeed, Figure 2.14 shows that treated incumbents are somewhat more likely to enter self-employment following an automation event, although the effect is very small – 0.3 percentage points cumulated over five years (a 6 percent increase relative to a five-year probability of 5.2 percent among the control group). This means self-employment is unlikely to be an important compensating income source.

Lastly, we find evidence that automation also leads to impacts on early retirement, defined as the receipt of retirement benefits prior to reaching the legal retirement age, as shown in Figure 2.14. In particular, five years after the automation event, treated incumbent workers are 0.3 percentage points more likely to be observed in early retirement. While these effects are similar in absolute size to those for self-employment, the average five-year probability of early retirement among control-group incumbents is much lower, around 1.7 percent. As such, the treatment effect represents an 18 percent increase in the incidence of early retirement.[33]

[33]However, early retirement does not account for all wage income losses: when we estimate our models separately for workers under the age of 55, we still find statistically significant relative wage income losses of the same size as in our full sample.

Taken together, these results show that automation-impacted workers are more likely to switch industries, but do not move to firms with different characteristics in terms of size, wages, or automation intensity. However, the documented income losses experienced by these workers are only partially offset by benefit systems and other income sources, implying automation-affected workers largely bear the adjustment cost themselves.

Figure 2.14: Cumulative probability of entering self-employment or early retirement

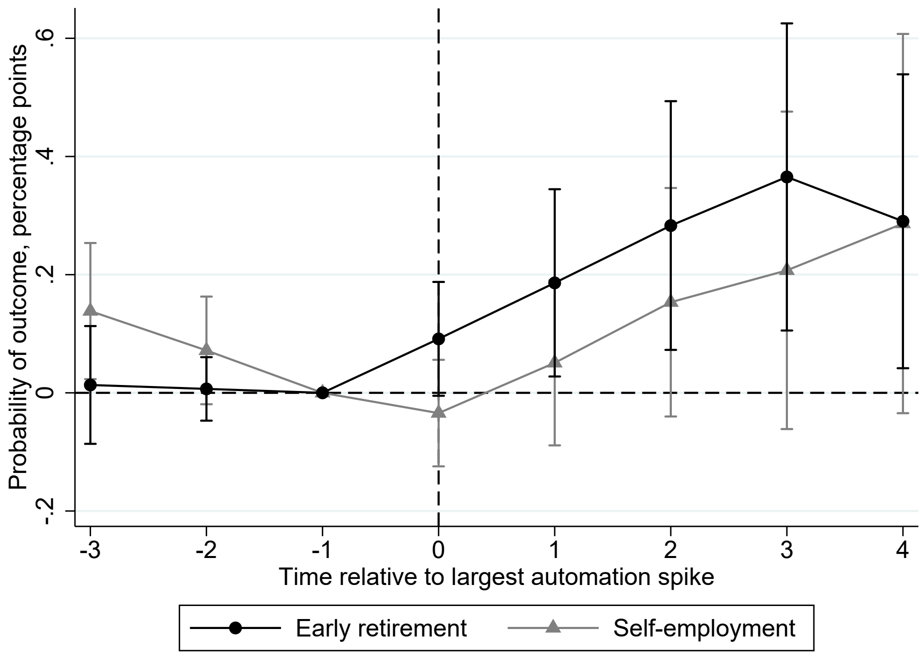

Notes: N=8,375,960. Whiskers represent 95 percent confidence intervals.

2.4.3 Robustness checks

Having laid out our main results, we show that these findings are robust to a number of alternative model specifications and other checks.

First, we subject our results to a randomization test as first introduced by Fisher (1935).[34] To do this, we take our sample of 36,490 firms, randomly draw firms with replacement, and then for each of these firms randomly assign a year to have a placebo automation event.[35] We then construct treated and control firms based on these placebo events. We repeat this procedure 100 times, where each permutation sample contains the same number of treated and control firms we have in our actual estimation sample.

Results are shown in Figure 2.15: each gray line presents a set of placebo (dynamic) treatment estimates, whereas the black line presents our actual treatment estimates. The graph also shows probability values calculated using the rank of the

[34] Also see Kennedy (1995) for an overview and Young (2018) for a recent application and evaluation of the value of these tests.

[35] Note that this permutates both the assignment of treatment to firms, and their timing across years, since both are part of our empirical procedure.

absolute value of our estimated coefficient among the 100 permutated estimates.[36] Something at least as extreme as our treatment estimate is unlikely to occur by chance: from the first post-event year onwards, the probability is very close to zero. Permutation estimates for two of our other three outcome variables also reject the null hypothesis: the hazard of leaving the firm and days in non-employment.[37] For the log daily wage impact, on the other hand, the randomization test shows that our point estimate is likely to occur when assigning automation spikes at random – this is as expected, since we did not find a statistically significant impact of the automation event on daily wages conditional on employment. All in all, this increases confidence that our estimates are not a statistical false positive.

However, even if our results are not occurring by chance, they may be driven by other real firm-level events that correlate with automation. Such events may impact labor demand at the firm level and thereby affect individual incumbent workers. Note that in our baseline specification, we do not see any pre-trends at the worker level, but the parallel trends assumption may of course still fail in the post-treatment period if such events coincide with the automation spike. We address this concern in three main ways.

First, we match additionally workers on their firms' pre-treatment employment trends. This implies we now ensure that treated and control workers are not only employed at firms that experience an automation event at some point in time, but where pre-treatment employment growth is similar. Second, our data include administrative information on some of these events, namely mergers, take-overs, acquisitions, firm splits, and restructuring.[38] As a second robustness check, we eliminate firms that experience such events anywhere in the estimation window. Third, we remove outlier firms in terms of employment changes (those experiencing an employment change exceeding 90 percent in any one year), both in the estimation window and outside of it.[39] The removal of these outliers is intended to capture any firm-level events which are not formally documented in our administrative records. Last, we remove firms where there was a new worker among the firm's top-decile annual wage income earners[40] in the three years prior to the automation event. This

[36]Results are very similar when using t-statistics rather than coefficient estimates to calculate probability values.

[37]See Appendix 2.7.3.

[38]We additionally observe firm births and deaths, but these are already excluded since we consider a balanced sample of firms over the observation window: we do allow firm births in the first year of observation, however.

[39]Results are similar when only removing firms with outliers inside the estimation window.

[40]Conditional on this worker earning at least 150 real euros a day, i.e. 40.000 euros a year.

is intended to capture automation events coinciding with managerial change, which may bring changes in personnel policy unrelated to automation.

Figure 2.16 summarizes the results for all three robustness checks pertaining to firm-level events (along with baseline model estimates). Estimates are very similar across the board, though effects are somewhat smaller when eliminating firms with (suspected) management change: this suggests that automation may sometimes be the result of a new manager changing business practices. Overall, however, our findings are very robust, showing that firm-level events other than automation are unlikely to be the driving force behind the worker impacts we find.

We have performed further robustness checks.[41] Specifically, we change our empirical specification by removing individual fixed effects (and replacing them with worker- and firm-level control variables), or additionally matching incumbent workers on firm size and on firm tenure (i.e. within the three year minimum firm incumbency requirement). We also consider various alternative spike definitions – including spikes in automation per worker rather than in total costs, and when measuring average costs only in the pre-event period. This all leads to very similar results. Lastly, we show that results are robust to varying the spike threshold from two to four times the average automation costs (our baseline is thrice the average automation costs). Estimated effect sizes are somewhat larger for higher compared to lower thresholds, as expected, but these differences are not statistically significant. This shows that our results are not driven by the specific spike size cut-off we employ in our baseline estimates.

[41] These are reported in Appendix 2.7.3.

Figure 2.15: A randomization test for relative wage income estimates

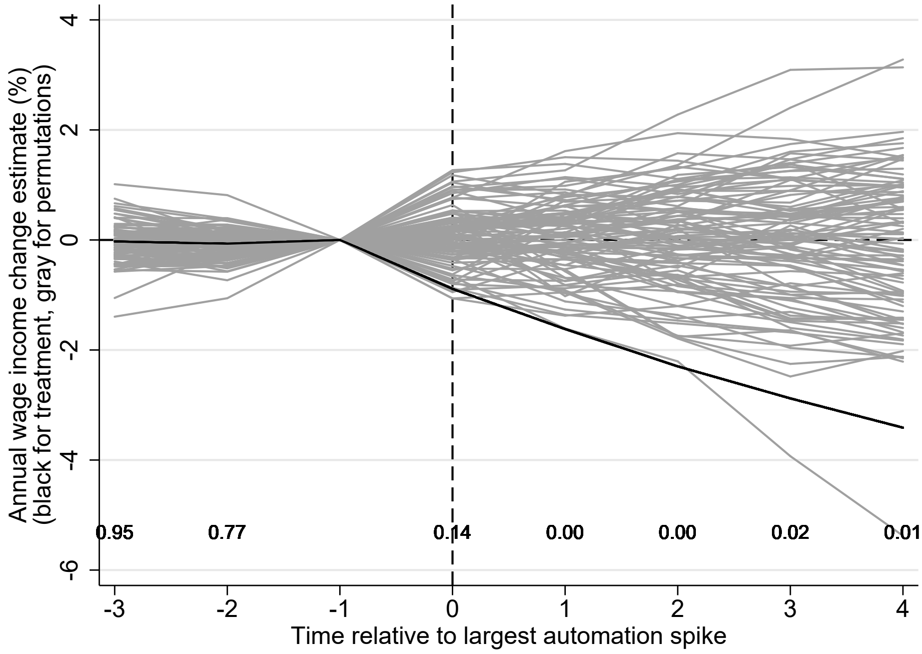

Notes: 100 permutations. The numbers printed at the bottom of the graph are probability values for the treatment estimates, based on the randomization test.

Figure 2.16: Robustness to removing other firm-level events

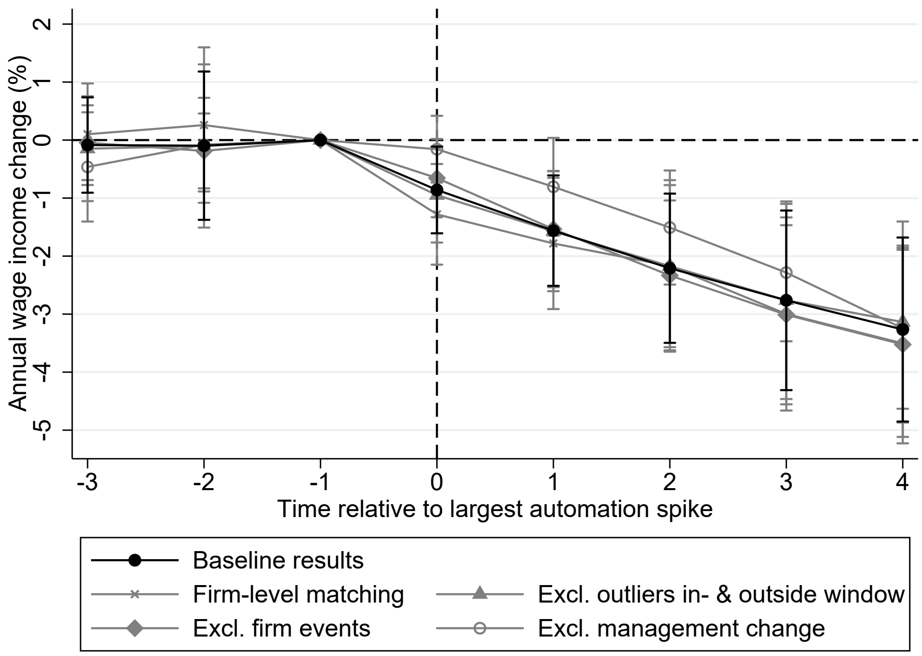

2.4.4 Effect heterogeneity

So far, we have shown income losses for incumbent workers driven by non-employment spells, only partially compensated by benefit payments. Here, we investigate further how different types of workers are affected: in particular, we consider effect hetereogeneity by incumbent age, gender, firm size, and sector of employment; and also separately study workers with less firm tenure than incumbents. Although our data lack a direct skill measure, we consider how impacts differ by age-specific wage quartiles, both overall and within firms. For succinctness, we will only show estimates for relative annual wage income, as this is the summary measure capturing all other impacts. Any noteworthy differences in results for firm separation, non-employment duration, log daily wages, and early retirement are described where relevant.

2.4.4.1 Incumbent worker characteristics

Here, we consider how incumbent workers with different characteristics fare after an automation event. For each of the groups considered here, we contrast the effect against the same group at the control firm by using an interaction term – this results in a decomposition of the mean effects found in section 2.4.1. In particular, we estimate the following model:

$$y_{ijt} = \alpha + \beta D_i + \gamma post_t + \delta_0 \times treat_i \times post_{it} + \sum_k [\delta_k \times treat_i \times post_t \times z_{ki}] + \lambda X_{ijt} + \varepsilon_{ijt}, \tag{2.3}$$

where, as before, i indexes workers, j firms, and t time relative to the automation spike. For succinctness, we estimate the average annual effect over the entire post-treatment period rather than reporting the year-by-year coefficients. As such, $post_t$ is a dummy variable indicating the post-treatment period (i.e. $t \geq 0$). Further, z_{ki} is a dimension of worker heterogeneity, such as gender, age in the year before automation, or age-specific wage rank, containing $k+1$ categories. In addition to the controls included in equation 2.2, X_{ijt} also contains z_{ki} as well as the interaction terms $z_{ki} \times treat_i$ and $z_{ki} \times post_i$. In equation 2.3, δ_0 gives the estimated treatment effect for the reference group, and δ_k the deviation from that effect for category k of worker characteristic z_i. βD_i capture worker fixed effects, and standard errors are clustered at the treatment level as before.

Table 2.7 summarizes how average post-treatment effects for annual wage income differ across workers of different ages, gender, and their (initial) firms' sector and employment size. First, we find that workers over the age of 50 are most negatively affected by automation events: while differences with younger age groups are not always statistically significant, the point estimates suggest all other groups experience somewhat smaller income losses. This is not because older workers leave the automating firm at higher rates, but rather, because they experience larger increases in non-employment duration. Unsurprisingly (and not reported here), the early retirement effects we found are entirely driven by the oldest workers. Taken together, older workers appear to face higher adjustment costs from automation than do younger ones.

We do not find any statistically significant differences in impact by gender and firm size. However, we may be underpowered in detecting differences across these groups. It should also be noted that effect heterogeneity across firm size (reported in column 3) is important for our purposes because our automation cost survey overrepresents large firms – while these of course employ the majority of workers, it could still bias the found worker-level effect of automation events by including too low a number of workers experiencing such events in small firms. For this reason, it is reassuring that displacement effects are found across the firm size distribution. If anything, losses are somewhat higher (though not statistically significantly so) for workers employed at smaller firms, which implies we would probably find somewhat higher average wage losses from automation if our data were more representative in terms of firm size. Lastly, although not reported here, we find that firm separation increases as a result of automation across all firm sizes, but most strongly so for the largest firms – the fact that this does not translate to larger wage losses for these workers suggests they have better outside options.

In column 4, we consider to what extent the impacts of automation differ depending on which sector the worker's firm belongs to: that is, our treatment effect is interacted with workers' sector of employment in $t = -1$. For this model, Manufacturing is the reference category. Note that sectoral differences may exist for various reasons. First, automation technologies may be sector-specific, and differ in terms of how much they displace labor. For example, it is possible that industrial or warehouse robots are more labor-replacing than automated check-out systems. Second, the workers employed in these different industries may have different characteristics (including unobservable ones), making the impacts differ. Third, to the extent that skills are industry-specific, sectoral labor market conditions matter: displacement would be more costly in sectors with an excess supply of workers. While we cannot distinguish between these different explanations, it is still important to consider whether our results are driven by displacement effects in a subset of sectors, or whether the found impacts are pervasive.

Our finding here is that automation leads to wage income losses that are quite pervasive across sectors: this highlights that robotics is likely not the only automation technology displacing workers from their jobs. The exception is Accommodation and food serving, where no income losses (nor increases in firm separation) are detected. However, Accommodation and food serving is also a sector with one of the lowest automation expenditures per worker (and the lowest number of automation events, see Table 2.5), as well as contributing only 2 percent of the sample of incumbent workers.[42] On the other hand, incumbent workers in Wholesale and retail and Manufacturing do experience earnings losses – together, these two sectors employ almost half of all incumbents in our sample (26 and 23 percent, respectively). We find that automation leads to increased firm separation rates for all sectors except Accommodation and food serving and Construction.[43] All in all, we find that automation events originating in different sectors have qualitatively similar impacts on workers.

Unfortunately, our data do not contain any occupation information, and only contain education information for a small and selected subsample of workers. Instead, we obtain a measure of workers' skill level by calculating each worker's wage rank by age in $t = -1$. We then group workers into quartiles based on this rank. For example, the top-quartile workers in this measure are those who earn in the top 25 percent of earnings across the sample for workers of their age in the year before the automation event.[44]

Results are reported in the first column of Table 2.8: workers in the highest age-specific wage quartile are used as the reference category. We do not detect any statistically significant differences: that is, workers across all wage quartiles experience displacement from automation. Indeed, the point estimates for deviations from effects for the top quartile are quite small.

The similarity of losses across the wage distribution may of course be partially driven by differences in the firms where automation spikes occur: lower losses for one "skill" group may be offset by higher exposure to automation events in our sample. While the estimates in column 1 matter for the average worker's exposure to displacement from automation, we are also interested in which workers are displaced

[42]See Table 2.18 in the Appendix.

[43]In a separate analysis, we find that the overall differences found between incumbents and recent hires in the next section are not only due to their different sectoral or firm size affiliations.

[44]As an alternative skill measure, we calculate residual wage quartiles (by first regressing worker wages in $t = -1$ onto a set of observables and their interactions): results (not reported here) are very similar.

Table 2.7: Relative wage income effects by incumbents' characteristics

(1) Age		(3) Gender	
Age 50+ (ref)	-3.04***	Male (ref)	-1.52***
	(1.15)		(0.57)
Deviations from reference group for:		*Deviations from reference group for:*	
Age <30	1.20	Female	-1.39
	(3.94)		(0.97)
Age 30-39	0.96	(4) Sector	
	(0.93)	Manufacturing (ref)	-1.98**
Age 40-49	1.61*		(0.99)
	(0.92)	*Deviations from reference group for:*	
(2) Firm size		Construction	1.05
500+ employees (ref)	-1.53		(1.73)
	(1.35)	Wholesale & retail trade	-2.23
Deviations from reference group for:			(1.51)
200-499 employees	1.21	Transportation & storage	0.71
	(1.77)		(1.79)
100-199 employees	-2.19	Accommodation & food serving	4.57**
	(1.77)		(2.32)
50-99 employees	0.17	Information and communication	-0.25
	(1.57)		(1.76)
20-49 employees	-2.18	Prof'l, scientific, & techn'l act's	-0.24
	(1.46)		(1.80)
1-19 employees	-2.06	Administrative & support act's	1.55
	(1.52)		(2.01)

Notes: Estimates from four separate models, N=8,375,960 for each model. All coefficients are average annual effects over the post-treatment period ($t = 0$ through $t = 4$); coefficients have been multiplied by 100. *p<0.10, **p<0.05, ***p<0.01.

within firms. Therefore, the second column in Table 2.8 reports estimates by workers' age-specific *within-firm* wage quartile. That is, the bottom quartile reflects incumbents who are in the lowest 25 percent of their firm's wage distribution for their age.[45] If anything, this reveals that the highest-paid workers by age *within* firms appear to lose more wage income than do lower quartiles, although these differences are not statistically significant.

We should be careful about drawing strong conclusions from these results since they may be capturing other factors than pure worker skill, such as the quality of the worker-firm match. However, Table 2.8 does provide an important insight that is counter to a common intuition: there is no evidence that workers lower down the wage rank for their age ("lower-skilled" workers) are displaced more often by automation

[45] Note that these quartiles cannot be calculated for the smallest firms: however, all previous findings are very similar in this subsample, suggesting that this is not driving the results.

events.[46] Since our approach captures a wide range of automation technology, this could be consistent with Webb (2019), who uses patent data to show that while low-skilled workers are most exposed to robotics, other automation technologies such as software and artificial intelligence impact more on work performed by medium- and high-skilled workers.

2.4.4.2 Worker tenure: incumbents versus recent hires

Our identification strategy for the impacts of automation is to consider individual workers who have a pre-existing working relationship with the firm, as evidenced by at least three years of firm tenure. We now turn to estimating our models for a second group of workers: those with less than three years of firm tenure prior to the automation event. Compared to incumbent workers, these workers have been hired relatively recently – we therefore refer to them as recent hires. This worker group is more likely to hold temporary contracts, which could imply different treatment effects. However, causal identification of the treatment effect for recent hires could prove more difficult as they may have been hired in anticipation of the automation event. We therefore analyze them separately, and generally put more stock in our results for incumbent workers.

On average, recent hires earn lower wages and spend a higher number of days in non-employment compared to incumbents.[47] They also have higher benefit receipts, and are more likely to be female, and younger. Compared to incumbents, recent hires are overrepresented in larger firms, and are most commonly employed in firms in Administrative and support activities, whereas incumbent workers are most often observed in the Manufacturing sector.

We now estimate equation 2.2 for recent hires in the same way we have for incumbents, while additionally creating a zero income bin when matching on pre-event income, and matching individual workers on pre-event trends in non-employment duration.[48] After matching, our sample contains $404,796$ unique recent hires $(78,282$

[46] These findings are confirmed when we estimate our models for the (small and selected) subset of observations where education level data is available.

[47] Table 2.17 in the Appendix shows summary statistics both incumbent and recently hired workers, showing averages and standard deviations across the balanced panel of workers and years.

[48] In particular, we estimate a linear trend in non-employment duration for individual recent hires before treatment, and match treated and control group recent hires using four bins of this trend: up to the 10th percentile, the 10th percentile to the median, the median to the 90th percentile and higher than the 90th percentile. Together with the other matching variables, we obtain 82,942 strata for recent hires, and can match 95 percent of treated recent hires (using 65 percent of control group recent hires).

Table 2.8: Relative wage income effects by incumbents' wage quartile

(1) Overall age-specific wage quartile		(2) Within-firm age-specific wage quartile	
Top quartile (ref)	-2.17**	Top quartile (ref)	-2.43*
	(1.06)		(1.27)
Deviations from reference group for:		*Deviations from reference group for:*	
Third quartile	0.39	Third quartile	0.79
	(0.84)		(0.86)
Second quartile	0.09	Second quartile	0.33
	(1.07)		(1.10)
Bottom quartile	-0.09	Bottom quartile	1.69
	(1.65)		(2.20)

Notes: The two models are estimated separately. 8,375,960 observations for column (1); 5,894,240 observations for column (2). All coefficients are average annual effects over the post-treatment period ($t = 0$ through $t = 4$); coefficients have been multiplied by 100. *p<0.10, **p<0.05, ***p<0.01.

of whom are treated): given our observation window of 8 years ($t = -3$ through $t = 4$) this results in $3,238,368$ observations.

Unlike for incumbent workers, we find no income losses from automation for recent hires, as shown in Figure 2.17. Relative to recent hires in the control group, point estimates are positive[49] but never statistically significant – hence, recent hires do not have different annual wage earnings as a result of automation. This could be the case because recent hires have built up less firm-specific human capital, and therefore are more able to adapt to new job tasks either within the same firm or when moving to a new employer. This is consistent with Carneiro et al. (2015) and Lefranc (2003), who find that income decreases following displacement result mostly from the loss of returns to accumulated firm-specific human capital. However, it may also be the case that recent hires do not lose income because these workers are in part hired in anticipation of the automation event – in this case their outcomes are endogenous to the event. Consistent with new hires being better matched (or able to adjust) to their firms' new technologies, we do not find any statistically significant increase in firm separation for these workers, and differences in non-employment duration with the control group are very close to zero over the entire pre- and post-treatment period.

Taken together, these results in this section show that although we detect some effect heterogeneity, our findings for incumbents are not driven by workers in a

[49]The point estimate suggests recent hires gain 4.5 percent of an annual income in total over five years following an automation event (corresponding to around 1,032 euros from a pre-treatment average income of 22,944 euros, see Table 2.18 in the Appendix).

Figure 2.17: Relative annual wage income effects for incumbents versus recent hires

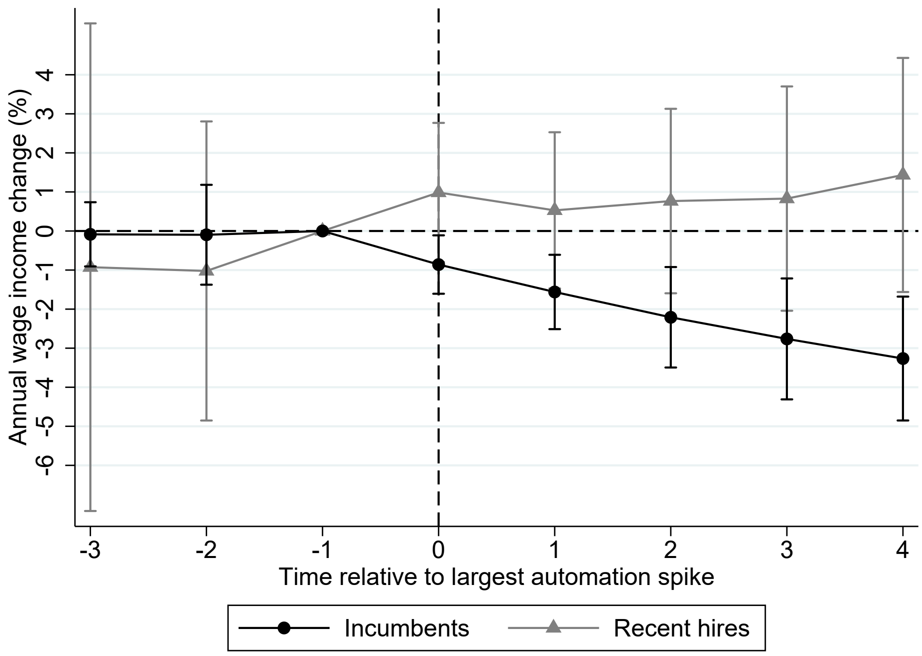

Notes: N=3,238,368 for recent hires and N=8,375,960 for incumbents. Whiskers represent 95 percent confidence intervals.

small subset of sectors, firm sizes, or age groups. Further, the income losses found for incumbent workers are not seen among recent hires, and automation affects incumbent workers from all ranks of the “skill” distribution.

2.5 Comparison to computerization

We have found that automation displaces incumbent workers: this raises the question whether this effect is specific to automation technologies or occurs with investment in new technology more generally. This question is also relevant from the perspective of recent theoretical frameworks which distinguish labor-replacing technologies from labor-augmenting ones (e.g. see Acemoglu and Autor 2011; Acemoglu and Restrepo 2018d): automation is typically modeled as being labor-replacing in nature.

While we do not have a specific technology inventory at the firm level, Statistics Netherlands conducts a separate and partially overlapping firm survey on investments, including computer investments.[50] This item is called ‘computers’ or ‘computers and

[50]Investments in software and in communication equipment are only measured from 2012 onwards, so we only consider computer investments. In 2012, software investments are of a similar magnitude as computer investments.

other hardware' (depending on the year) and consistently defined as follows: "All data-processing electronic equipment insofar as they can be freely programmed by the user, including all supporting appliances. Do not include software.". All investment within the company counts towards the expenditures, also if the equipment is second-hand, or leased or rented, or produced within the company. It excludes investments in plants that are located abroad or resulting from take-overs of other organizations whose operations are continued without change.

In this section, we analyze the effects of computer investments in a similar way to that of automation investment, and directly contrast it to the impacts of automation in the part of the sample where we have overlapping data. This serves two purposes. First, as outlined above, we can consider to what extent spikes in automation costs have different effects on workers than do spikes in computer investment. Second, since automation cost and computer investments are somewhat correlated at the firm level, we can remove firms which have computer investment spikes within our estimation window to rule out that our automation event is partially capturing investment in computers. Conversely, we will also estimate the effects of computer investments in isolation, that is, excluding any events where automation spikes occur within the estimation window.

We first show some summary statistics on computer investments (section 2.5.1) before turning to the comparison between automation and computerization (section 2.5.2). Throughout, we consider the overlapping sample of firms where we observe both automation cost and computer investment data. This means our dataset has a smaller number of observations, and is more skewed towards larger firms as these are most likely to be sampled in both surveys. However, our results are qualitatively identical in the full samples for both automation and computerization.

2.5.1 Summary statistics

Here, we show summary statistics on both computer investments and automation costs for the overlapping sample of firms where we observe both: this allows for the most direct comparison.

Table 2.9 informs on the distribution of automation costs and computer investment across firms and years. Automation costs are higher than computer investments across the distribution, both in total and per worker. Of course, it should be noted that both can come with other unmeasured correlated costs, such as software for computers, and machinery for automation.

Further, Tables 2.10 and 2.11 compare automation and computer investments across firms of different sizes and sectors. In both tables, we also report the ratio of observed automation to computer expenditures per worker. As expected, Information and communication has the highest computer investment per worker, followed by Professional, scientific, and technical activities. Accomodation and food serving and Construction have the lowest computer investment per worker. When considering the relative importance of automation and computer technology, Manufacturing is the most automation-intense compared to other sectors, whereas Information and communication is the most computer-intense. These patterns are reassuring. Like for automation, we generally see higher computer investment per worker for larger than smaller firms, but the pattern is less dramatic: this is reflected by the ratio of automation to computer expenditures rising with firm size.

Lastly, Figure 2.18 plots quantiles of the distribution of computer investments per worker over time. Unlike for automation costs, investments per worker are declining initially (perhaps reflecting the aftermath of the dot-com bubble) and relatively flat thereafter. Of course, effective computing power per worker is likely to have grown: computer investments have been deflated by the overall price index, which is unlikely to capture quality improvements in computing equipment.

2.5.2 Automation versus computerization

In order to compare automation to computerization, we construct computer investment events in the same way we have for automation, but using computer investment per worker.[51] We use the same threshold, assigning firms a computer investment spike if their computer investment per worker exceeds three times their usual level.

The resulting distribution of computerization events is reported in Table 2.12. Compared to automation events, computerization events are more frequent. However, Figure 2.19 shows that computerization events are also clearly visible in the estimation sample: in the event year, treated firms spend around 1,605 euros per worker, compared to around 246 euros in the years before and after. This is similar to the 1,912 euros treated firms spend per worker during automation events.

Armed with our overlapping sample and both types of events at the firm level, we now construct four different datasets. First, we consider automation events and

[51]This is necessary because, unlike automation expenditures, computer investments are not part of total costs; and because total investments are inconsistently defined over our sample period. In Appendix 2.7.3.5 we also define automation spikes based on outlays per worker (rather than based on cost shares), and find very similarly sized effects.

Table 2.9: Computer and automation cost share distributions

	Automation cost		**Computer investment**	
	level	*per worker*	*level*	*per worker*
p5	0	0	0	0
p10	0	0	0	0
p25	0	0	0	0
p50	16,747	297	5,554	99
p75	69,617	957	31,042	447
p90	241,274	2,175	112,889	1,126
p95	568,915	3,518	250,652	1,868
mean	249,275	1,032	99,666	559
mean excl. zeros	346,396	1,434	155,619	873
N firms × yrs	171,549		171,549	
N firms × yrs with 0 costs	48,098		61,680	

Notes: All numbers are in 2010 euros. The number of observations is the number of firms times the number of years.

Table 2.10: Automation costs and computer investments by sector

Sector	**Autom. cost per worker**	**Comp. inv. per worker**	**Ratio autom. to comp.**	**Nr of obs** *Firms*	*Firms × yrs*
Manufacturing	998	369	2.7	5,153	40,773
Construction	497	215	2.3	2,821	18,319
Wholesale & retail trade	1,152	544	2.1	7,220	50,381
Transportation & storage	917	456	2.0	2,279	15,834
Accommodation & food serving	256	151	1.7	742	4,462
Information & communication	2,030	2,420	0.8	1,562	9,756
Prof'l, scientific, & techn'l activities	1,272	772	1.6	2,345	14,708
Admin & support activities	863	388	2.2	2,914	17,316

Notes: Overlapping sample.

Table 2.11: Automation costs and computer investments by firm size

Firm size	**Autom. cost per worker**	**Comp. inv. per worker**	**Ratio autom. to comp.**	**Nr of obs** *Firms*	*Firms × yrs*
1-19 employees	2,233	1,091	2.0	2,267	11,326
20-49 employees	851	543	1.6	10,451	66,339
50-99 employees	838	456	1.8	5,804	41,460
100-199 employees	944	500	1.9	3,418	26,466
200-499 employees	1,204	569	2.1	1,929	16,202
≥500 employees	1,640	637	2.6	1,167	9,756

Notes: Overlapping sample.

computer events in isolation: that is, we identify treated and control group workers for one type of event while ignoring the other. This allows us to estimate our DiD

Figure 2.18: Firm-level computer investment per worker over time

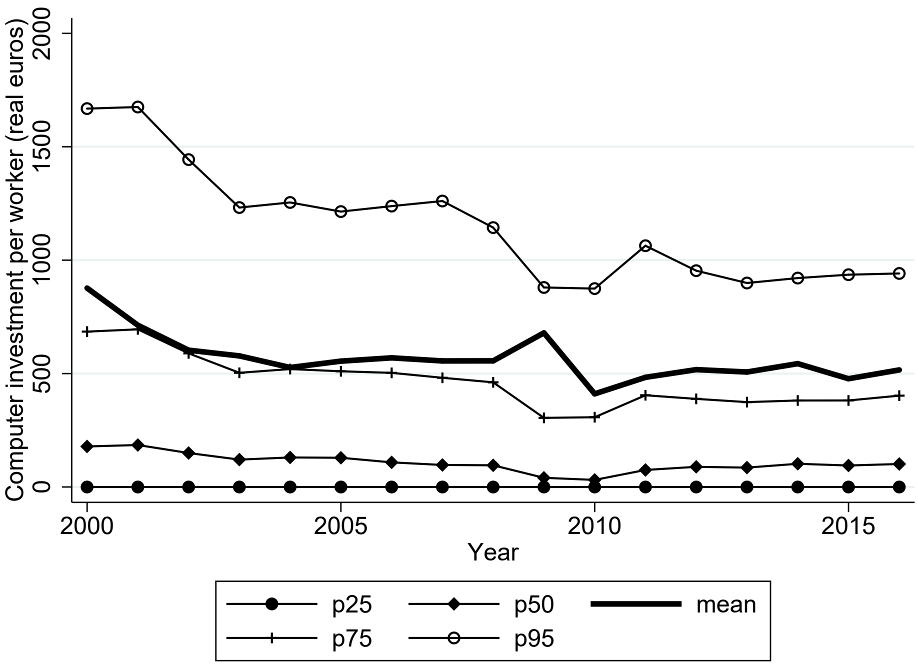

Table 2.12: Automation costs and computer investments by firm size

	Percentage of firms with event type:	
Nr of events	Automation	Computerization
0	71.8	47.9
1	22.5	41.9
2	4.8	9.1
3	0.7	1.1
4	0.1	0.1

Notes: Overlapping sample, N=25,036 firms.

model for automation and computerization separately. However, these two events are correlated across firms over time: that is, firms that have recently had one type of event are more likely to also experience the other sometime soon – sometimes even in the same year. This implies any estimated impact of automation may be contaminated by computerization, and vice versa. We therefore construct two additional samples of events which occur in isolation: that is, we only retain those automation (computerization) events where there is no computerization (automation) event occurring in the estimation window for either treated or control group firms. For each of the four samples, we then estimate equation 2.2 and report results in Figure 2.20.

Figure 2.19: Computer investment per worker for treated firms

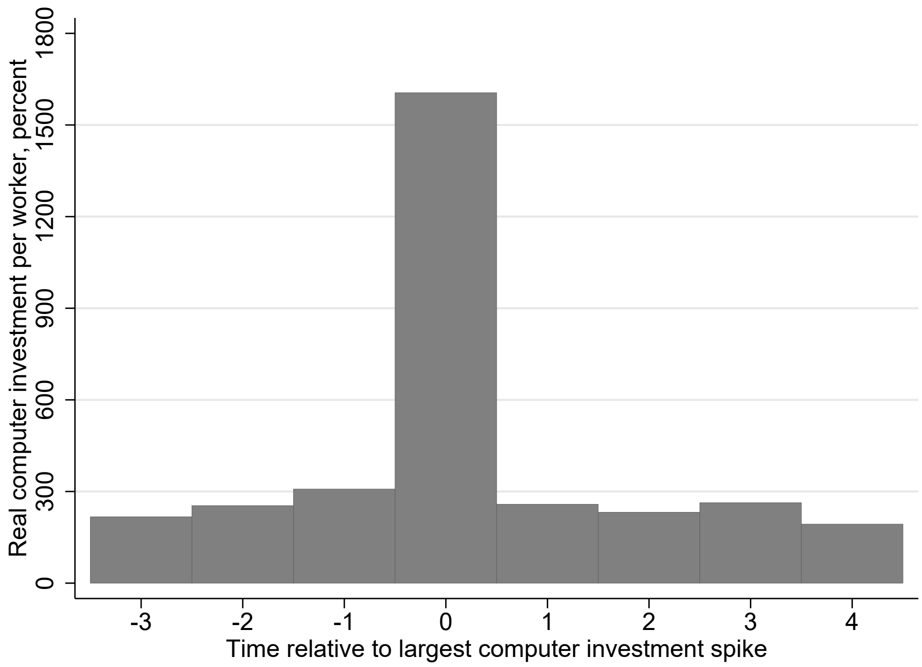

Notes: Overlapping sample, N=2,745.

This comparison leads to several findings. First and foremost, computerization does not lead to income losses for incumbent workers: estimates are small and never statistically significant. This is in contrast to automation, which does lead to income losses. Further, the income losses of automation are larger when removing concurrent computerization, and the effects of computerization on income are (slightly) smaller when removing concurrent automation. Consistent with these results, we do not find any increase in firm separation or non-employment duration for workers impacted by computerization. This may be because automation is generally more labor-displacing in nature than computerization, but also because computer technologies are already further along the adoption curve (as also evidenced by the higher frequency of computer investment spikes). This could imply that during computer spikes, it is mostly older vintages of computer capital that are being replaced, while any displacement of workers by computers has already played out in the past. All in all, irrespective of the reason, automation is (currently) a more labor-displacing force than computerization from the perspective of a firm's incumbent workers.

Figure 2.20: Relative annual wage income effects of automation and computerization

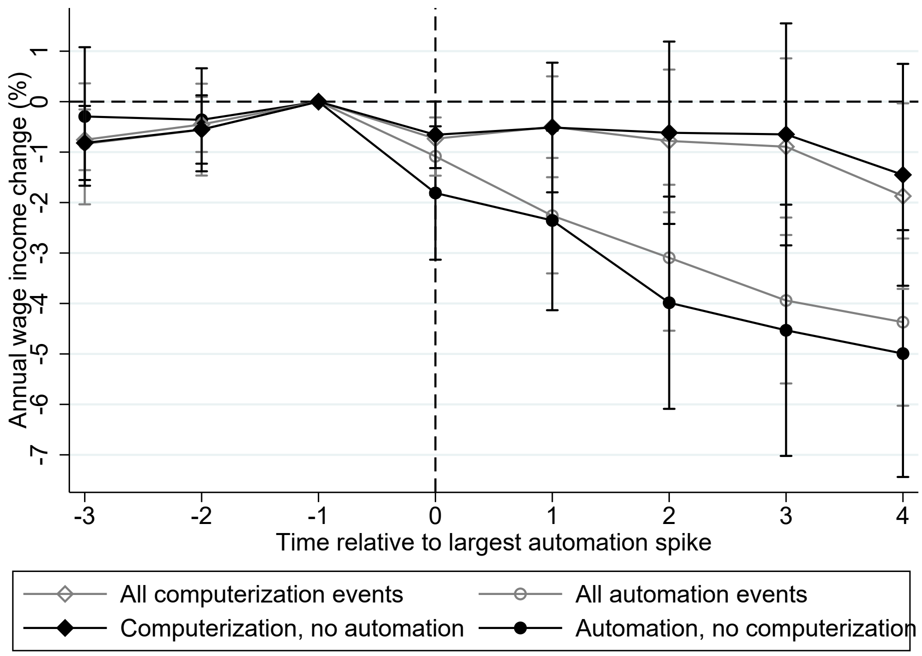

Notes: All estimates are for the overlapping sample where we observe data on both automation costs and computer investments. N=10,217,088 for all computerization; N=7,783,929 for all computerization excluding automation; N=8,110,456 for all automation; and N=4,632,880 for automation excluding computerization.

2.6 Conclusion

We provide the first estimate of the impacts of automation on individual workers, using firm-level data on automation expenditures across all non-financial private sectors in the Netherlands over 2000-2016. Leveraging a novel differences-in-differences design exploiting automation event timing, we show that automation at the firm significantly increases incumbent workers' hazard of separating from their employers. This firm separation is followed by a decrease in days worked, leading to a five-year cumulative wage income loss of some 11 percent of one year's earnings. Wage income losses are only partially offset by various benefits systems, and older workers are more likely to enter early retirement.

In contrast to displacement from mass lay-offs, however, workers do not experience daily wage scarring as a result of automation events, nor do we find evidence that automation displaces the firm's more recent hires. Further, automation events affect relatively few workers and these effects occur more gradually: two percent of incumbent workers separate from their employers in the first year, followed by a trickle of ongoing separations.

Our findings are robust to a range of specification and falsification tests, including controlling for other firm-level events such as mergers and acquisitions, firm restructurings, and suspected management change. Furthermore, effects are quite pervasive across different incumbent worker types, as well as firm sizes and sectors. Put simply, where work is automated, incumbent workers face a higher risk of displacement, resulting in adjustment costs. These adjustment costs may partly explain anxiety about new workplace technology expressed in opinion surveys, despite ample macro-economic evidence on technology's economic benefits. In contrast, we do not find evidence that workers face adjustment costs from firms' investments in computer technology. This suggests that, from the perspective of incumbent workers, automation is (currently) a more labor-displacing force.

Our findings of course do not imply that automation destroys jobs on net in the economy. As a related macro literature has shown, there are various countervailing mechanisms which our models do not inform about, including effects operating through firms' input-output linkages and changes in final demand. However, by focusing on workers directly impacted by automation events, our results contribute to understanding the adjustment costs of automation, which matter for science and policy alike.

2.7 Appendix

This supplemental appendix contains details on sample construction, additional summary statistics, and additional robustness checks on our baseline specifications.

2.7.1 Sample construction

We follow several steps to construct our sample. We start with 36,490 firms for which we observe at least 3 years of automation costs data in the Production Statistics (PS) survey for 2000-2016. Then for each firm we determine if they have at least one spike, with spikes defined as explained in section 2.3.1. We keep all firms with at least one spike observed over 2000-2016. This leaves us with 10,476 firms. We lose an additional 2 firms because we cannot merge them to administrative worker records, so we continue with 10,474 firms. Then, for each calendar year y we define a set of potential treatment and control group automation events as follows.

Potential treatment events for y are defined as a firm having its largest spike in y. y has to lie between 2003 and 2011, so that for each event we at least have a window of three years before and five years after the event. Events are excluded if the firm also has another spike in the $t = [-3, 4]$ window around the event. This gives us 2,446 potential treatment group events. Note that we do not require that firms are observed in the PS survey in all 8 years around the potential event, but they do have to exist in each year in the window. Figures 2.4 and 2.5 are based on this sample. Figures 2.6 and 2.7 furthermore require firms to be observed in the PS survey in all 17 years, that is from 2000 to 2016.

Potential control events for y are defined as firms that have their largest spike in year $y + 5$ or later. Hence, these spikes have to occur between 2008 and 2016. Furthermore, events are excluded if a firm has another spike in the $t = [-8, -1]$ window around the event. Again, we do not require that firms are observed in the PS survey in all the years surrounding the event, but the firm does have to exist in this period. This gives us 21,575 potential control events.

Columns (1) and (2) in Table 2.13 show the number of potential treatment and control events per calendar year. Note that our procedure implies that multiple control group events can involve the same firm, but for different years y. It is also possible that one treatment group event and one or more control group events involve the same firm in different years y. For example, a firm that has its largest spike in 2010 can be a potential treatment event in 2010, but also serve as a potential control event for treatment events in 2003, 2004, or 2005. Similarly, a firm having its largest spike in 2011 can serve as a control group event for treatment events in 2003, 2004, 2005, or 2006. For our 21,575 potential control events, 20,838 involve a firm that is involved in more than one potential control event, while 737 events involve a firm that is involved in only one potential control event. Firms with potential control events are on average involved in 6.3 potential control events, with a maximum of 9 events. For our 2,446 potential treated events, 1,021 involve a firm that is also involved in at least one potential control event in another year and 1,425 involve a firm that is not involved in a potential control event.

We then merge our firm-level data to worker data and keep only events for which we can find at least one incumbent worker who is between 18 and 65 years old at $t = -1$. This leaves us with 2,439 potential treatment events merged to 124,225 incumbent workers and 21,399 potential control events merged to 1,157,536 incumbent workers.

We then take these samples of potential treatment and control group events and incumbent workers and apply our matching procedure. The details of our matching procedure are discussed in the main text in section 2.3.4. During matching we also apply some basic sample selection procedures to remove outliers. In particular, we remove students, people with total wage earnings above 0.5 million euros in a year or 2,000 euros per day and people earning less than 1/4 of the fulltime minimum wage (5,000 euros per year) or less than 10 euros per day on average. We also drop people whose earnings or daily wages increase more than tenfold since the year before treatment. After these sample selection procedures and matching, we are left with 102,599 treated incumbents in 2,429 events and 944,396 control incumbents in 21,175 events. Columns (3) and (4) in Table 2.13 show the number of events in each calendar year after matching. This is our main balanced panel of 1,046,995 workers, observed for 8 years, which means we have 8,375,960 observations in total.

2.7.2 Additional descriptives

Here we present some additional summary statistics for firms' automation expenditures and events; and for the main estimation sample of workers.

2.7.2.1 Summary statistics on automation expenditures and events

First, Table 2.14 shows automation costs per worker across firms of different sizes, when outliers in automation costs have been removed – that is, observations with automation costs per worker in the 99th percentile or above are dropped. This

Table 2.13: Number of treatment and control events at the firm level by calendar year

	Potential events		Events after matching	
Calendar year	Control	Treatment	Control	Treatment
2003	3,492	199	3,411	199
2004	3,271	200	3,208	198
2005	2,965	200	2,922	198
2006	2,725	223	2,688	223
2007	2,452	321	2,392	320
2008	2,205	311	2,161	308
2009	1,926	363	1,894	361
2010	1,544	315	1,521	312
2011	995	314	978	310
Total	**21,575**	**2,446**	**21,175**	**2,429**
Unique firms involved	4,588	2,446	4,543	2,429
Unique firms only used once	737	2,446	751	2,429

Notes: Table show the number of potential treatment and control events, and the number of events remaining after matching, for each calendar year.

removes fewer than 50 firms, predominantly one-person firms with high automation outlays per worker.

Second, Figures 2.21 and 2.22 show automation cost spikes in both shares and levels per worker for a balanced sample of firms – that is, firms which are observed in the Production Statistics survey every single year over 2000-2016.

Third, Table 2.15 shows how spiking and non-spiking firms differ in terms of observables. In particular, it estimates a firm-level linear probability model where the dependent variable is a dummy for the firm having at least one automation spike over 2000-2016: this model is estimated for the sample of firms where we observe at least 3 years of automation cost data. This highlights that firms that we observe having automation spikes are different from those where we do not observe a spike. In particular, the smallest firms are less likely to experience an automation event, and there are some sectoral differences, with automation events least likely to be observed in Accommodation and food serving, and most likely in Information and communication.

Lastly, firm-level spike timing is not easy to predict based on observables. Specifically, a predictive model with observables performs only marginally better than a random prediction where we uniformly distribute spikes across years where the firms are observed. This is reflected in the Brier (1950) skill scores for ten k-folded samples reported in Table 2.16. These are constructed as follows. We draw a 10 percent random sample without replacement from the main sample of spiking firms, and do this ten times: these are the test samples. The remaining 90 percent of observations for each of these test samples constitute the ten training samples. We then estimate a logit model with firm fixed effects and time-varying observables[52] on each training sample and predict the probability of having a spike in a year for each test sample, assuming that each firm will have exactly one spike. We also calculate the spike probability by year per firm from random prediction, simply as one over the number of years the firm is observed. For the model-based and random predictions in each of the ten test samples, we calculate the Brier score, defined as the mean squared difference between the prediction and the actual outcome. Lastly, we obtain the Brier skill score as $1 - \frac{Brier_{model}}{Brier_{random}}$, reflecting the percent prediction improvement of the model relative to random prediction. This improvement is 2.9 to 4.2 percent, confirming that spike timing is hard to predict.

Figure 2.21: Automation cost share spikes for treated firms, balanced sample

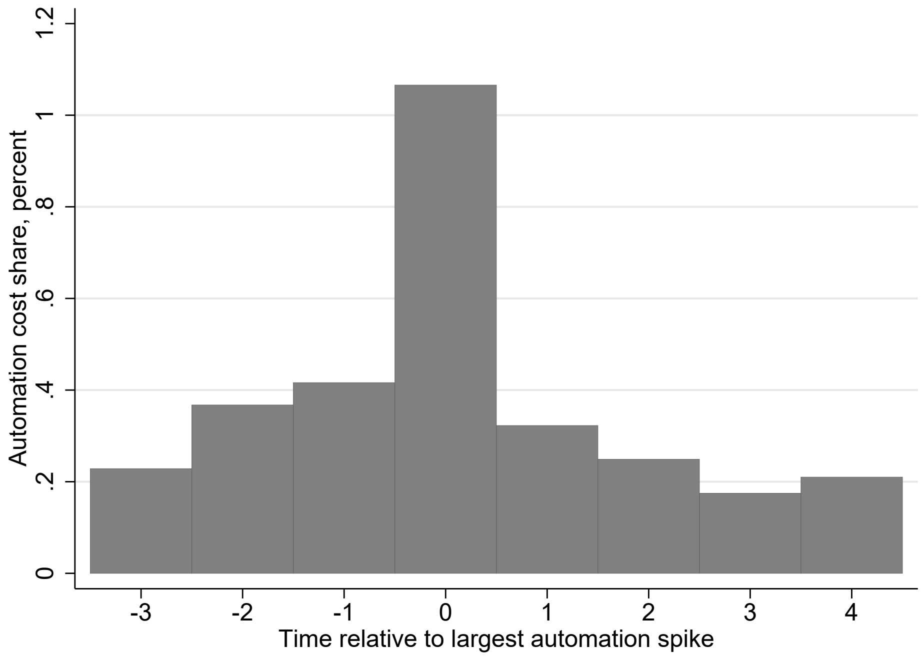

Notes: N=315 in all years.

[52]These observables are: firm average log yearly and daily wages, log total wage bill, log number of workers, log average worker age, log average worker tenure at the firm, share female and a full set of interactions. Results are similar when we do not include these interactions, or additionally include lagged observables in the model.

Table 2.14: Automation costs by firm size class after removing outliers

Firm size class	Total cost	Cost per worker		Cost share (%)		Nr of obs	
	Mean	Mean	SD	Mean	SD	*Firms*	*Firms × yrs*
1-19 employees	9,028	613	1,074	0.35	0.75	9,836	48,378
20-49 employees	20,393	639	1,068	0.37	0.74	13,755	86,523
50-99 employees	48,418	690	1,099	0.39	0.73	6,282	46,756
100-199 employees	105,035	754	1,172	0.41	0.70	3,470	28,472
200-499 employees	287,917	930	1,421	0.45	0.75	1,966	17,600
≥500 employees	1,584,027	968	1,522	0.62	0.99	1,134	10,188

Notes: Automation cost level in 2010 euros, automation cost shares as a percentage of total costs, excluding automation costs. Total N firms is 36,443; Total N firms × years is 237,917.

Figure 2.22: Automation cost level per worker for treated firms, balanced sample

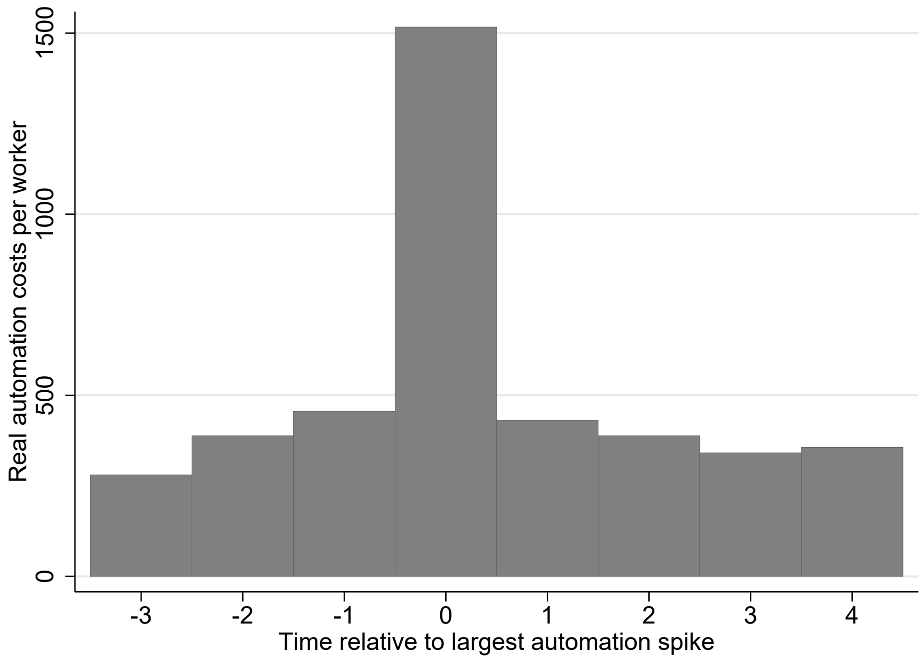

Notes: N=315 in all years.

2.7.2.2 Summary statistics for workers

Table 2.17 provides summary statistics on our sample of workers across all years, before matching. Table 2.18 provides summary statistics on our matched sample of workers. For both incumbents and recent hires, we show the averages and standard deviations for the dependent as well as independent variables used in our models, separately for the treated and control group. Note that we have $102,599 + 944,396 = 1,046,995$ observations for incumbents and $78,282 + 326,514 =$

Table 2.15: Correlates of a firm ever having an automation spike

Mean worker age	-0.0027***	Manufacturing	*reference*
	(0.0005)		
Share of women	-0.0075	Construction	-0.0294***
	(0.0126)		(0.0093)
Mean real annual wage / 1,000	0.0009***	Wholesale & retail trade	-0.0057
	(0.0002)		(0.0080)
1-19 employees	*reference*	Transportation & storage	0.0268***
			(0.0102)
20-49 employees	0.0347***	Accommodation & food serving	-0.0333**
	(0.0060)		(0.0151)
50-99 employees	0.0495***	Information & communication	0.0977***
	(0.0075)		(0.0114)
100-199 employees	0.0291***	Prof'l, scientific, & techn'l act's	0.0598***
	(0.0091)		(0.0102)
200-499 employees	0.0534***	Admin & support act	0.0089
	(0.0114)		(0.0099)
$\geq$500 employees	0.0250*	Constant	0.3304***
	(0.0142)		(0.0214)

Notes: 36,489 observations, each observation is a unique firm. The dependent variable is having an automation spike at any point in the sample. Standard errors in parentheses * p<0.1 ** p<0.05, *** p<0.01

Table 2.16: Brier skill scores for predicting automation spikes

Sample	**N**	**Brier skill score**
1	127,378	0.040
2	127,232	0.037
3	126,485	0.038
4	127,886	0.033
5	126,812	0.042
6	126,440	0.029
7	127,599	0.033
8	126,954	0.032
9	126,497	0.028
10	127,616	0.036

$404,796$ for recent hires: given our observation window of 8 years ($t = -3$ through $t = 4$) this adds up to the $1,179,584 \times 8 = 8,375,960$ observations for incumbents and $404,796 \times 8 = 3,238,368$ for recent hires used in our regressions.

Table 2.17: Descriptives for all workers

	Incumbents (1)	Recent hires (2)
Annual wage income	37329.65	25872.14
	(25159.74)	(22092.42)
Daily wage if employed	156.91	126.32
	(94.43)	(80.91)
Annual non-employment duration (in days)	28.59	66.32
	(61.96)	(90.25)
Hazard of leaving the firm	0.04	0.12
	(0.21)	(0.32)
Total benefits	403.45	1610.27
	(2724.93)	(4449.22)
Probability of entering early retirement	0.01	0.00
	(0.09)	(0.05)
Probability of becoming self-employed	0.03	0.04
	(0.18)	(0.21)
Share female	0.25	0.35
	(0.43)	(0.48)
Foreign born or foreign-born parents	0.16	0.27
	(0.36)	(0.44)
Age	42.60	36.60
	(10.24)	(10.06)
Calendar year	2006.90	2006.86
	(3.37)	(3.43)
Manufacturing	0.37	0.15
	(0.48)	(0.36)
Construction	0.11	0.07
	(0.32)	(0.26)
Wholesale and retail trade	0.19	0.16
	(0.39)	(0.36)
Transportation and storage	0.09	0.07
	(0.28)	(0.26)
Accommodation and food serving	0.02	0.02
	(0.13)	(0.15)
Information and communication	0.06	0.05
	(0.23)	(0.23)
Professional, scientific, and technical activities	0.08	0.08
	(0.28)	(0.28)
Administrative and support activities	0.09	0.38
	(0.29)	(0.49)
0-19 employees	0.06	0.06
	(0.23)	(0.23)
20-49 employees	0.14	0.12
	(0.35)	(0.32)
50-99 employees	0.12	0.10
	(0.32)	(0.30)
100-199 employees	0.12	0.10
	(0.33)	(0.30)
200-499 employees	0.15	0.11
	(0.36)	(0.31)
$\geq$500 employees	0.42	0.52
	(0.49)	(0.50)
Observations	9,017,448	4,392,416

Notes: Unweighted means for all worker-year observations. Standard deviations in parentheses.

Table 2.18: Descriptives on matched worker samples

	Incumbents		Recent hires	
	Treated	*Control*	*Treated*	*Control*
	(1)	(2)	(3)	(4)
Annual wage income	35885.04	35950.43	22944.41	22963.90
	(23717.08)	(23973.53)	(17594.43)	(17658.74)
Daily wage if employed	145.77	146.13	109.72	109.10
	(90.79)	(90.11)	(64.91)	(64.29)
Annual non-employment duration (in days)	21.51	22.48	64.09	63.84
	(43.22)	(44.25)	(71.02)	(71.60)
Total benefits	0.00	0.00	1401.83	1406.47
	(0.00)	(0.00)	(3648.65)	(3688.33)
Probability of entering early retirement	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)
Probability of becoming self-employed	0.03	0.03	0.04	0.03
	(0.16)	(0.17)	(0.18)	(0.18)
Share female	0.33	0.31	0.43	0.38
	(0.47)	(0.46)	(0.49)	(0.48)
Foreign born or foreign-born parents	0.18	0.17	0.31	0.33
	(0.38)	(0.38)	(0.46)	(0.47)
Age	40.78	40.65	36.01	35.27
	(10.18)	(10.05)	(10.15)	(10.03)
Calendar year	2006.41	2006.41	2007.09	2007.09
	(2.39)	(2.39)	(1.99)	(1.99)
Manufacturing	0.23	0.23	0.07	0.07
	(0.42)	(0.42)	(0.25)	(0.25)
Construction	0.08	0.08	0.04	0.04
	(0.28)	(0.28)	(0.19)	(0.19)
Wholesale and retail trade	0.26	0.26	0.12	0.12
	(0.44)	(0.44)	(0.32)	(0.32)
Transportation and storage	0.09	0.09	0.05	0.05
	(0.28)	(0.28)	(0.22)	(0.22)
Accommodation and food serving	0.02	0.02	0.02	0.02
	(0.15)	(0.15)	(0.12)	(0.12)
Information and communication	0.05	0.05	0.04	0.04
	(0.21)	(0.21)	(0.19)	(0.19)
Professional, scientific, and technical activities	0.11	0.11	0.07	0.07
	(0.31)	(0.31)	(0.25)	(0.25)
Administrative and support activities	0.17	0.17	0.60	0.60
	(0.38)	(0.38)	(0.49)	(0.49)
0-19 employees	0.07	0.06	0.04	0.04
	(0.25)	(0.24)	(0.19)	(0.20)
20-49 employees	0.16	0.16	0.09	0.10
	(0.37)	(0.36)	(0.29)	(0.30)
50-99 employees	0.12	0.13	0.07	0.09
	(0.33)	(0.33)	(0.25)	(0.29)
100-199 employees	0.12	0.13	0.07	0.09
	(0.32)	(0.34)	(0.26)	(0.28)
200-499 employees	0.14	0.15	0.07	0.10
	(0.35)	(0.36)	(0.25)	(0.30)
$\geq$500 employees	0.39	0.37	0.67	0.58
	(0.49)	(0.48)	(0.47)	(0.49)
Observations	102,599	944,396	78,282	326,514

Notes: Weighted means for the full regression sample at $t = -1$, where weights are obtained from coarsened exact matching as described in section 2.3.4. Standard deviations in parentheses.

2.7.3 Additional robustness checks

In this section, we present several additional robustness checks on our results. First, we present estimates for wage income in levels. Second, we show permutation estimates for each of our other dependent variables (the hazard of firm separation, days in non-employment, and log daily wages). Next, we change our empirical specification in a number of ways. Lastly, we consider how changing the spike definition changes our results.

2.7.3.1 Effects in levels

Figure 2.23 shows estimates for our main specification in levels. This shows that incumbents lost around 3,842 euros over five years in total, which is very similar to the 3,839 euros lost when estimating impacts on relative income (shown in Figure 2.8).

Figure 2.23: Annual real wage income in levels

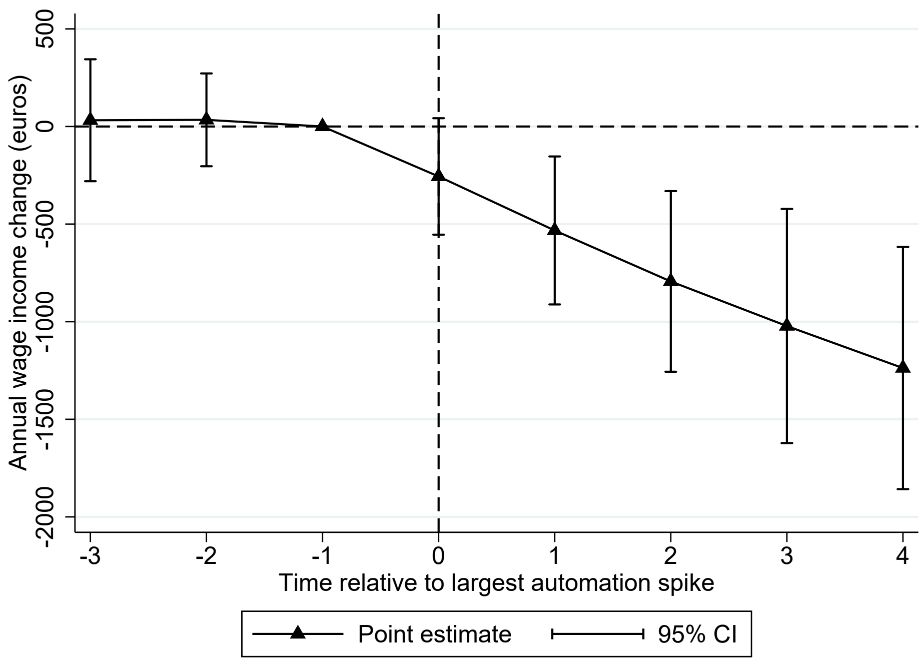

Notes: N=8,375,960. Whiskers represent 95 percent confidence intervals.

2.7.3.2 Randomization tests for other outcome variables

Figure 2.24 shows permutation estimates for the hazard of firm separation (panel a), days in non-employment (panel b), and log daily wages (panel c). Two-sided probability values are reported below each estimate: see the main text for a discussion of these results.

Figure 2.24: Additional randomization tests

(a) Probability of firm separation

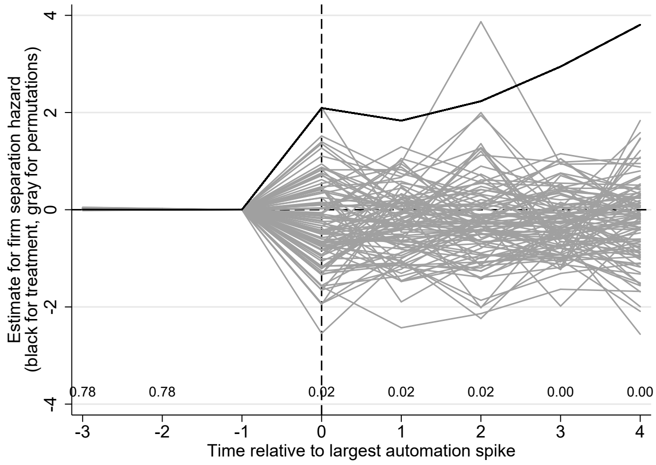

(b) Days in non-employment

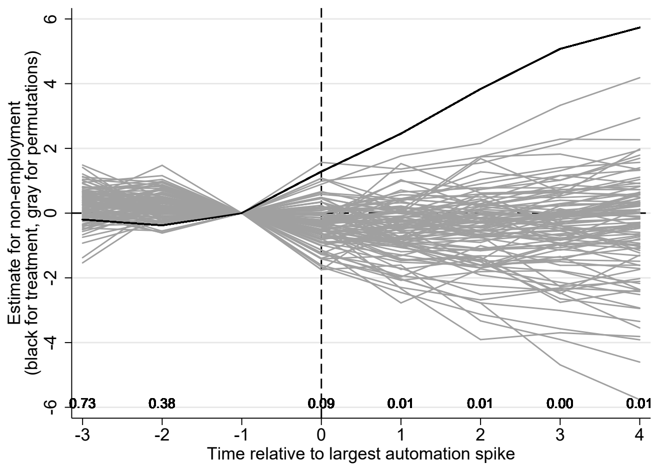

(c) Log daily wage

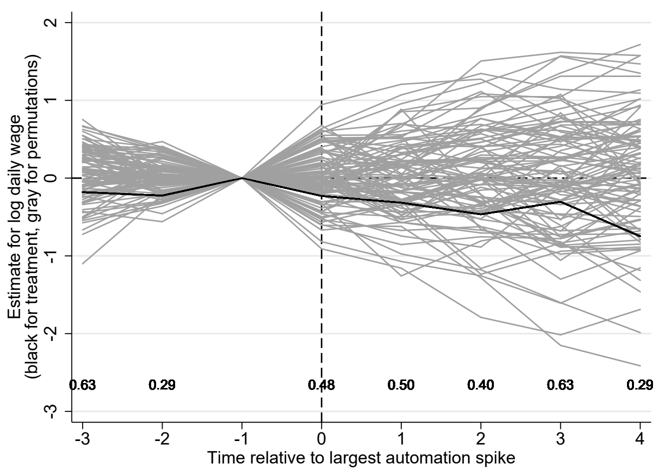

Notes: 100 permutations. The numbers printed at the bottom of the graph are probability values for the treatment estimates, based on the randomization test.

2.7.3.3 Changes in model specification

Here, we change our model specification in a number of ways. In particular, compared to our baseline estimates, Figure 2.25 shows results when additionally matching workers on their firm tenure in years (that is, beyond the three years of firm tenure that all treated and control group workers have); additionally matching workers on firm size; and when removing individual fixed effects from the model (these are then replaced by dummies for worker gender and nationality, as well as fixed effects for firm size categories, and for firm sector). Although estimates without individual fixed effects are a little less precise, results are extremely robust to these changes in specification.

Figure 2.25: Robustness to changes in model specification

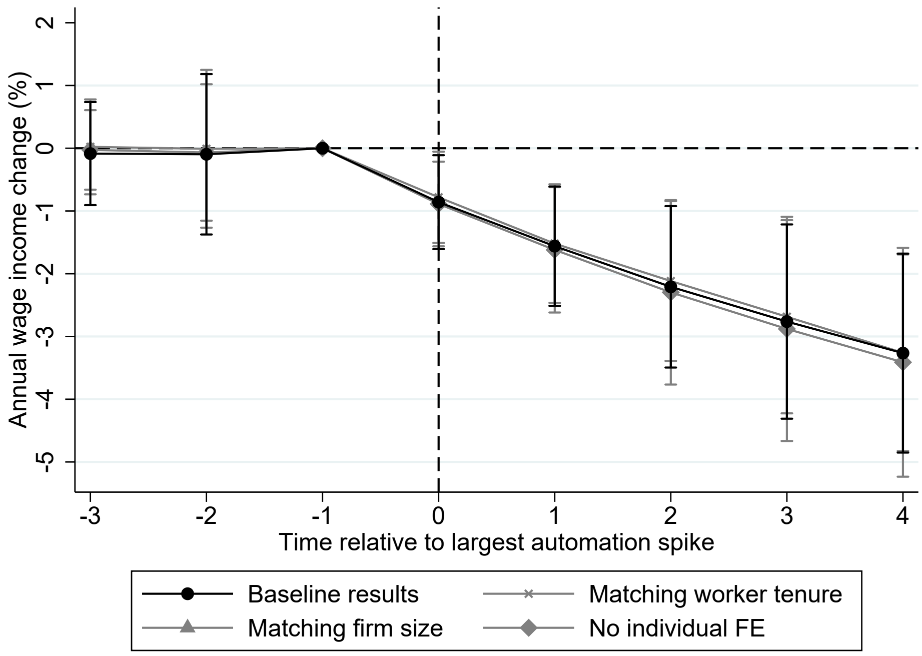

2.7.3.4 Removing other firm events

Figure 2.26 shows that our results for firm separation are also robust to excluding other firm-level events: see section 2.4.3 in the main text for a description.

Figure 2.26: Robustness to removing other firm events

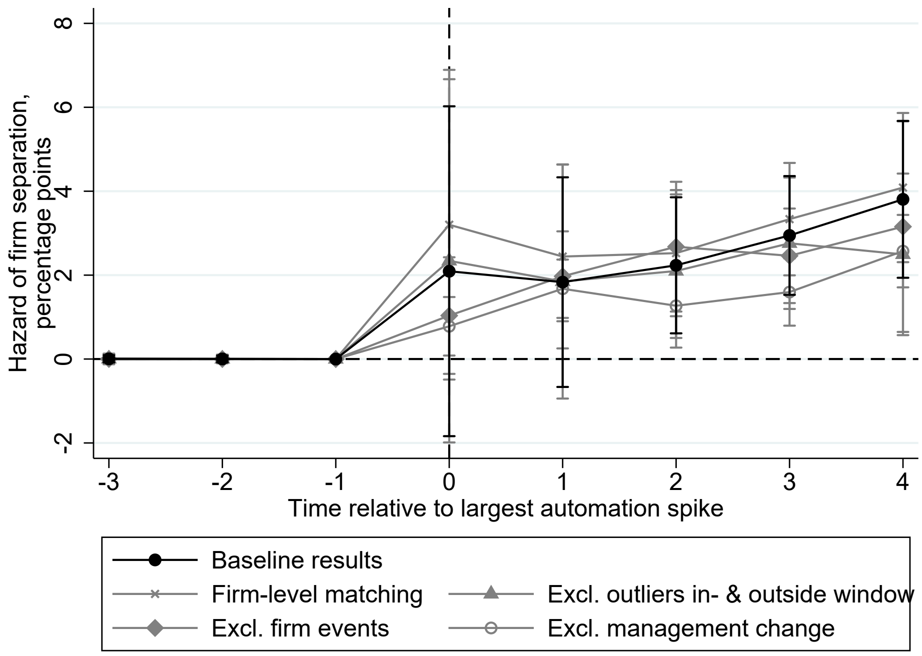

2.7.3.5 Alternative automation spike definitions

We have considered a range of alternative ways of identifying automation spikes, either by changing the spike definition or the spike threshold.

In particular, rather than using automation cost shares (i.e. automation costs in total costs), we can construct automation events from sharp increases in automation outlays per worker. This is more in the spirit of a literature studying the impact of increasing the number of robots per work. Within this event definition, we then also vary the point(s) in time where we measure employment (i.e. the denominator in the spike variable) – either for the years where we have data on total costs ("AC/worker"); or for the full set of years ("AC/worker, full emp data"); or only for the years pre-dating the candidate automation event ("AC/worker, pre-event emp data"). All variations produce similar results to our baseline estimates, as seen in panel (a) of Figure 2.27.

Further, we show that results are robust to varying the spike threshold from two to four times the average automation costs (our baseline is thrice the average automation costs). Panel (b) in Figure 2.27 reveals that estimated effect sizes are somewhat larger the higher the threshold, as expected, but these differences are not statistically significant. This highlights that our results are not driven by the specific spike size cut-off we employ in our baseline estimates.

Figure 2.27: Robustness to changes in spike definition

(a) Changes in spike definition

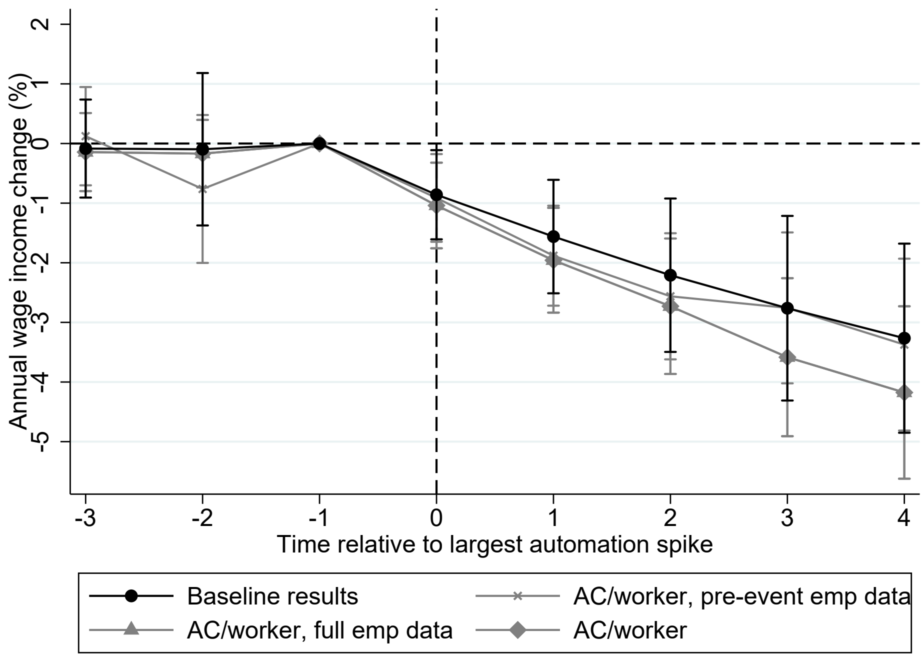

(b) Changes in spike threshold

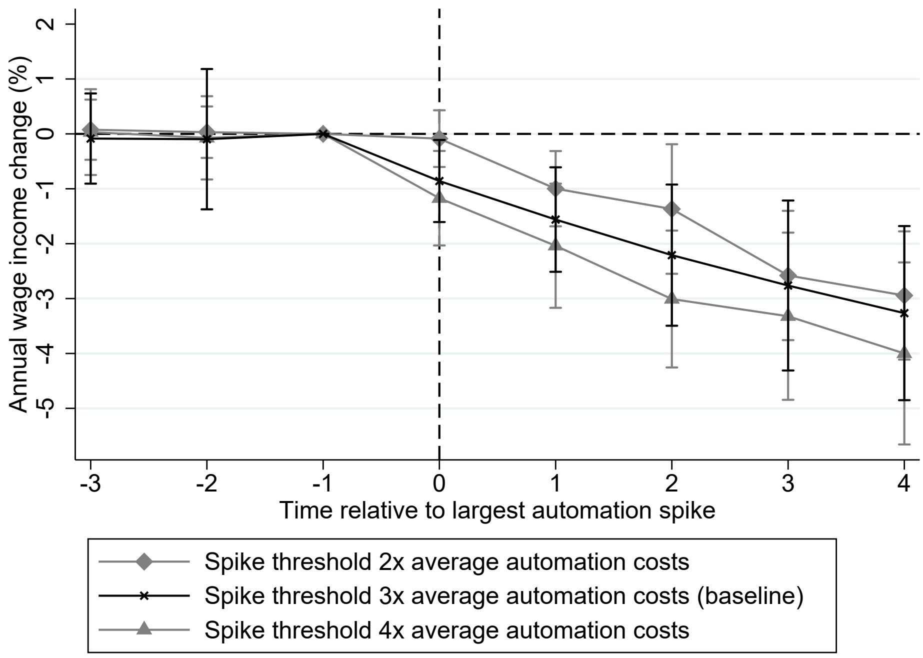

CHAPTER 3

Bad Start, Bad Match? The Early Career Effects of Graduating in a Recession for Vocational and Academic Graduates*

3.1 Introduction

Youth unemployment is a cause for concern in many countries. Especially during the Great Recession, when youth unemployment rates rose quickly in many OECD countries, there have been widespread worries about unemployment disrupting young people's lives and giving them a false start on the labor market. While short-term negative effects of entering the labor market in a recession are to be expected, some worry that young people will suffer long-lasting negative effects. If true, this type of hysteresis could lead to a lost generation of young workers who will be stuck in mismatches and low-paying jobs.

In this paper I consider the effects of labor market conditions at the moment of graduation on the early careers of tertiary educated graduates in the Netherlands.[1] Throughout the paper I examine the effects separately for graduates from universities, who take an academic track, and graduates from universities of applied science, who

*This chapter has been published as Van den Berge (2018).

[1] I examine graduates with degrees in ISCED categories 6 and 7. See section 3.2.3 for more detail.

take a more vocationally oriented track. In terms of flexibility, the Dutch labor market is somewhere in between the very flexible labor markets in the US and the UK and the stringent labor markets of many other European countries. I use administrative matched employer-employee data on graduates from 1996 to 2012. My data allow me to follow graduates on the labor market for up to eight years after graduation. I employ a measure of field-specific employment conditions at graduation. This best approximates the labor market conditions that high-educated graduates face.[2]

The paper addresses two questions. First, does graduating in a recession affect graduates from the academic and vocational tracks in higher education differently? I consider effects on wages, but also on employment, self-employment and whether graduates are dependent on benefits. Second, how do vocational and academic graduates catch up to their peers who started in good times? I look at the quality of firms workers start at and consider how they recover through job mobility by looking at whether workers climb the job ladder and estimating the wage returns to job mobility.

With respect to the first question I find that academic graduates suffer strong inititial wage effects of 10% for each percentage point decline (around half of a standard deviation) in field-specific employment at graduation. The wage losses gradually decline until they fade out after about five years on the labor market. The initial wage losses for vocational graduates are significantly smaller at close to 6% for each percentage point decline in field-specific employment at graduation. They remain significantly smaller than for university graduates in the first four years. However, wage losses for vocational graduates remain persistent at about -1% up to at least 8 years after graduation. Employment probabilities for both academic and vocational graduates are negatively affected in the first three to four years on the labor market. While self-employment is not affected for vocational graduates, for academic graduates I find evidence of graduates substituting regular employment for self-employment in the first years after graduation. This is not persistent though, as for later years I find evidence for a reverse substitution of self-employment with regular employment. Finally, I find virtually no effects on benefit take-up.

On the second question I find that job mobility plays a critical role in recovering from initial wage losses for both academic and vocational graduates who start in a

[2]This measure is similar to Beiler (2017). It uses the employment conditions in the sectors that students with a given field of study usually end up in as the measure of labor market conditions at graduation.

recession. They are both more likely to switch firms and sectors, and when they do switch, they gain more than their counterparts who started in a boom. Graduates are more likely to start in firms that pay lower wages in a recession and gradually move to higher paying firms. Both are also more likely to be mismatched in their early career. Interestingly, while switching sectors solves the initial mismatch for academic graduates, vocational graduates remain in sectors that are not typical for their field of study. This could at least partly explain the persistent wage losses for vocational graduates.

This study relates to two strands of literature. First, it relates to the literature on the effects of bad starting conditions at graduation on long-term labor market outcomes. Several papers have found that people who enter the labor market during a recession indeed suffer lower wages up to ten years or longer (Kahn, 2010; Oreopoulos et al., 2012; Brunner and Kuhn, 2014). This suggests that hysteresis might be a real problem, although more recent papers, covering both Europe and the US, find smaller losses for university educated graduates that disappear after three to five years on the labor market (Hershbein, 2012; Altonji et al., 2016; Liu et al., 2016; Cockx and Ghirelli, 2016; Speer, 2016). While most papers on the US only find lower wages for labor market entrants who started in a recession, papers looking at less flexible labor markets such as Belgium, Spain or Japan also find higher probabilities of non-employment (Genda et al., 2010; Cockx and Ghirelli, 2016; Fernández-Kranz and Rodríguez-Planas, 2018). The literature also finds that high-educated workers generally suffer more in terms of wages, while low-educated workers suffer more in terms of employment probabilities (Genda et al., 2010; Cockx and Ghirelli, 2016; Speer, 2016).[3]

The paper improves on this literature primarily by considering the effects separately for academic and vocational graduates. While some papers have explored differences between high and low educated workers (e.g. Oyer, 2006, 2008; Genda et al., 2010; Speer, 2016; Cockx and Ghirelli, 2016) there is as far as I'm aware no

[3]There is also a small Dutch literature on this topic. Van Ours (2009) and Fouarge (2009) show that there are no long-term differences in unemployment rates between cohorts entering the labor market in the recession of the 1980s and cohorts who entered just before. Erpelinck and Van Sonsbeek (2012) look at multiple cohorts and find that primarily tertiary educated graduates from the early 90s suffer long-term negative effects on their wages. Wolbers (2014) uses repeated cross-sections covering 1993 - 2011 and finds only short-term negative effects of entering the labor market during a recession on employment and job level. Limitations of this literature are that they use cross-section data, that they do not know the actual moment of labor market entry and that they do not take into account possible selection bias due to people adjusting their moment of labor market entry to the labor market conditions.

study that examines the differential effects for academic (general) and vocational graduates.[4]

The paper also adds to this literature by considering in detail how job mobility contributes to catching up. Other papers also consider the mechanisms of catching up. Motivated by a model of task-specific human capital (as in Gibbons and Waldman, 2004) most papers consider the role of the first firm in explaining the initial and persistent losses (Oreopoulos et al., 2012; Brunner and Kuhn, 2014; Liu et al., 2016). They generally find that the first employer plays an important role in explaining the losses. Recent papers also highlight the role of mismatch. They find that workers who start in a recesison are more likely to work in sectors that are not typical for their field of study (Liu et al., 2016; Altonji et al., 2016). Oreopoulos et al. (2012) also consider the role of job mobility for Canadian college graduates. They find that job mobility increases in the first five years on the labor market. Primarily graduates at the top end of the skill distribution are more likely to switch firms, while those at the bottom remain stuck at lower quality firms. After the first five years, the remaining catching up is within the firm. In addition to considering job mobility and firm quality, this paper is the first to consider how young workers recover through climbing the job ladder. There is recent evidence that during recessions the quality of vacancies is lower and that the probability of moving up the job ladder declines (Moscarini and Postel-Vinay, 2016; Haltiwanger et al., 2017). This paper confirms for the Netherlands that workers are indeed more likely to start at lower rungs of the job ladder in a recession. While most workers recover through job mobility, I do find some evidence that up to 8 years after graduation workers are more likely to remain lower on the job ladder.

Finally, my data allow me to take into account selection bias much more thoroughly than most other papers. A causal effect of the unemployment rate at graduation on later outcomes is only identified if students do not adjust their timing of graduation to labor market conditions. I find evidence that students do seem to adjust their timing. They are more likely to obtain an additional degree if economic conditions at graduation are bad, and the composition of the graduation cohort is different in bad economic conditions. I employ an IV strategy that requires

[4]A somewhat related study by Humburg et al. (2017) does examine the effect of the unemployment rate at graduation and self-reported “field-specific” skills on the probability of being unemployed or over-educated. They use a survey on European graduates for 17 countries. They find that graduates with high field-specific skills are less likely to be unemployed 5 years after graduation, but they find no interaction between field-specific skills and the unemployment rate at graduation. However, they also don’t distinguish between vocational and academic tertiary educated workers.

detailed data on when students enter higher education and their expected duration to deal with this problem. Other studies often do not observe the date of entry or graduation, and impute the year of school leaving using date of birth and expected school duration.[5]

Second, the paper relates to the literature on the benefits of vocational versus general education. Most of this literature has focused on whether vocational education eases the transition from eduation to the labor market.[6] The findings are mixed. Some studies indeed find that students with vocational education have better employment outcomes in their early career (Hanushek et al., 2017) but others find no differences (Fersterer et al., 2008). Studies focusing on the long-run outcomes generally find no differences between general and vocational tracks (Oosterbeek and Webbink, 2007; Malamud and Pop-Eleches, 2010; Hall, 2016). I primarily add to this literature by considering the effects of bad starting conditions on the transition from education to the labor market for vocational and general educated workers.[7] If vocational education eases the transition from education to work, one might expect that it also helps with finding a job in a recession. On the other hand, if there are fewer vacancies, and especially high-quality vacancies, it could be that general skills are more helpful in finding a job. Another contribution is that this paper looks at graduates from tertiary education, while the literature is mostly focused on graduates from (upper-)secondary education.[8]

The paper proceeds as follows. In section 3.2 I present my empirical strategy and discuss the data. Main results are presented in section 3.3. In section 3.4 I explore the mechanisms behind the initial wage losses and catching up and section 5.7 concludes.

3.2 Empirical strategy and data

I aim to estimate the effects of economic conditions at graduation on later labor market outcomes of vocational and academic graduates. In this section I first discuss

[5]Exceptions are Oreopoulos et al. (2012), who present some estimates using the same IV strategy as I do, but finds no differences, and Kondo (2015), who predicts age of entry by the highest degree attained, rather than using the actual date of entry into education.

[6]See Ryan (2001) and Wolter and Ryan (2011) for surveys on this literature.

[7]Hall (2016) is the only paper who considers the effect of the unemployment rate at graduation in a robustness analysis and she only examines it for workers who already have at least 10 years of working experience. She finds no differences between those with more general and those with more vocational education.

[8]Some exceptions are Heijke et al. (2003); Verhaest and Baert (2015); Humburg et al. (2017).

how I calculate a measure to approximate the economic conditions at graduation. Then I present my empirical model. Finally, I will describe the data, discuss the construction of important variables and present descriptive statistics.

3.2.1 Economic conditions measure and model

To best approximate the economic conditions at graduation, I exploit the fact that the field of study students graduate in provides them with the skills that typically match to a given set of industries (Liu et al., 2016). For example, students graduating in finance generally have skills that are well suited for working in the financial sector compared to students graduating in healthcare. As a result, finance graduates have probably been hit harder in the financial crisis than students graduating in healthcare. To arrive at a measure of economic conditions for each field of study, I follow Beiler (2017). I first calculate for each field of study the share of workers holding a degree in that field in each industry. These shares indicate the importance of each industry for a field of study. Then I calculate the year-on-year percentage change in employment for each industry, and use the field of study specific shares as weights to arrive at a weighted year-on-year employment change for each field of study. The measure is calculated as follows

$$e_{cf} = \sum_{s} w_f^s * \Delta e_{sc} \tag{3.1}$$

where s indicates sectors of industry, f field of study and c the year of graduation (cohort). The variable Δe_{sc} is defined as $\frac{e_{sc}-e_{sc-1}}{e_{sc-1}}$ and denotes the year-on-year (from year $c-1$ to year c) percentage change in employment in each sector. The time-invariant weights for each industry within each field of study are given by w_f^s.[9] The sample used to construct the measure is discussed in section 3.2.5, where I also present descriptive statistics.[10]

[9] The weights sum to one within each field of study.

[10] In an earlier version of this paper I used a different measure for economic conditions. I calculated the unemployment rate for all workers with a degree from a field of study using microdata from the Labour Force Survey. However, this measure is subject to measurement error and endogeneity issues. Measurement error arises from small sample sizes due to a small number of unemployed workers in the Labour Force Survey for some fields of study. The current measure is not subject to this type of measurement error since it uses all workers employed in a sector. Endogeneity results, as pointed out by a referee, from workers who could be unemployed because they graduated in an *earlier* recession. If the effects of graduating in a recession are indeed persistent, this could bias the unemployment measure. This effect also exists in the current measure, but it is much smaller, because graduates make up only a small part of the total number of workers in each sector.

I then use the weighted employment changes for each field of study to estimate the effects of the economic conditions at graduation on labor market outcomes. I use the following linear model

$$Y_{itcf} = \alpha + \beta_{exp} exp_{it} * e_{cf} + \zeta X_i + \delta_{exp} + \phi_c + \mu_f + \tau_t + \varepsilon_{it}, \tag{3.2}$$

where Y is the outcome variable (wage, employment status or some other labor market outcome) for individual i observed in year t who graduated in cohort c in field of study f. I control for a full set of potential experience fixed effects δ_{exp} (with potential experience exp defined as years since graduation), cohort fixed effects ϕ, calendar year fixed effects τ and field of study fixed effects μ. X_i is a vector of time-constant individual control variables: age at graduation and gender. The coefficients of interest are the β_{exp}'s which describe the change in the experience profiles caused by a one percentage point change in field-specific employment e_{cf} at graduation. I allow the effect to differ for each year of potential experience.[11] For example, β_0 describes the effect of a one percentage point change in field-specific employment at graduation in the year of graduation, while β_1 describes the same effect in the first year after graduation. I estimate the effect for the first 8 years after students obtain a degree, so for $exp = 1$ until $exp = 8$.[12]

I take the year of graduation of a student's highest degree as their point of entry into the labor market. The potential experience fixed effects pick up the average effect of potential experience on the outcome variable. The year fixed effects control for any variation in labor market conditions and for other year effects that might affect wages apart from the change in employment at graduation or experience. The cohort fixed effects pick up changes at the cohort level that might affect labor market outcomes, such as the increased participation rate in higher education or changes in student support.[13] Finally, field of study fixed effects control for average differences in the labor market opportunities of students with different fields of study. To take into account that individuals from the same cohort might have experienced similar

[11] I have also experimented with more restricted functions for experience, such as a quadratic or cubic. The results are similar to the more flexible version I use here.

[12] I observe some cohorts (from 1999 until 2007) more than 8 years, while other cohorts (2009 - 2012 and 1996 - 1998) are observed less than 8 years on the labor market. I always observe at least 4 years for each cohort. Extending the analysis to 10 years yields similar results and does not affect the main conclusions.

[13] Since cohort, potential experience and year fixed effects can not be identified at the same time, I have to impose another restriction. I follow the literature and impose that one additional year effect is zero (Oreopoulos et al., 2012; Cockx and Ghirelli, 2016).

shocks - e.g. changes in the education system - I cluster standard errors at the level of the graduation cohort and field of study.[14]

3.2.2 Selection bias

OLS estimates of β_{exp} will be biased if the field-specific change in employment at graduation is correlated with unobserved variables that also determine the outcome variable. A potential source of selection bias is that students adjust their timing of graduation to labor market conditions. Students could postpone graduating during a recession or students could leave school earlier during a boom if they already found a job. Students could also pursue an additional degree. To deal with this potential source of selection bias I employ an IV strategy similar to Oreopoulos et al. (2012).[15]

As an instrument for the change in the employment rate at graduation I use the change in the employment rate at the predicted year of graduation using the actual entry date and the nominal study duration. This is a valid instrument if it is not related to labor market outcomes, except through the unemployment rate in the actual year of graduation. The instrument would be invalid if, at entry, students are able to predict the state of the labor market at their expected graduation date. Given the difficulty of predicting unemployment rates, this seems a plausible exclusion restriction.

Another way to think of the difference between the OLS and IV strategy is the following. Ideally we would like to compare two people who are identical, except that one graduates during a recession and the other one does not. We are effectively looking for random allocation of recessions at graduation. This is of course impossible, so we have to resort to comparing people who graduate at different points in time. We can then choose to compare two groups of people. One, people who might have started at different points in time, but graduate at the same point in time. In this case, variation in labor market conditions at graduation arises from differences in study duration and differences in the moment people enter higher education. Two, people who start at different points in time, but study for the same amount of time. Estimates using the first group could be biased, because they might have

[14] Clustering at the graduation cohort yields similar results.

[15] In section A.1 in the Appendix I present evidence that students indeed seem to adjust to labor market conditions. This confirms the need for this IV strategy. I find that graduates from a university of applied science are more likely to pursue a university degree if conditions are bad when they obtain their initial degree. University graduates are also more likely to pursue an additional degree. Finally, graduates from a university of applied science are somewhat older and less likely to be female if starting conditions are bad.

adjusted their timing of graduation to labor market conditions. On the other hand, estimates using the second group only exploits variation in labor market conditions at graduation due to differences in the year of entry into higher education. OLS estimation uses both groups to estimate the effects, while my IV strategy effectively only uses the second group.[16]

I operationalize the IV strategy as follows. I predict the year of graduation using the nominal duration of a particular study. In this setting a 'study' is defined as a completed program in a particular field, such as a bachelor's in engineering or a master's in philosophy. If a student takes both a bachelor's and a master's, the total nominal duration of her study is the sum of the nominal duration of the respective bachelor's and master's program. The nominal duration is based on the number of ECTS required, where one year is equal to 60 ECTS. For many degrees it is 4 years, but for some degrees (e.g. medical or technical academic degrees) the nominal duration is 6 years. I then use the change in employment in the predicted year of graduation as an instrument for the change in employment in the actual year of graduation.

3.2.3 The Dutch higher education system

In my analysis I use data on graduates from the Netherlands. Since the Dutch higher education system differs somewhat from the US and other countries, I will discuss the institutional setup in some detail. At the start of secondary education, Dutch students are tracked in three levels. Only the highest two tracks give direct access to higher education. Tracking is based on a standardized test taken at the end of primary school and primary teacher's evaluations.

The second track (HAVO) takes five years and gives direct access to the vocationally oriented universities of applied science (*hoger beroepsonderwijs*, the literal translation would be "higher vocational education"). These are similar to for example the *Fachhochschule* in Germany.[17] The highest track in secondary education

[16]Even if there is no selection in study duration, there could still be selection due to cohort effects. An illustration of such remaining cohort effects could be the choice of field of study at the start of the higher education career. Before entering higher education, students choose a field of study. Their choice might be influenced by the labor market conditions at the moment of choosing their field. For example, if students make their choice of field during a recession, they might be more likely to pick a field with more secure or higher labor market returns than during a boom. It is beyond the scope of this paper to examine in detail how labor market conditions at the end of secondary school affect student's choice for post-secondary education. I include field of study and cohort fixed effects to take this into account as much as possible.

[17]Students in the lowest track in secondary education have the opportunity to go to a university of applied science if they finish their vocational degree (MBO) first. This takes a total of seven or eight years. I exclude these students from my analysis, so I will not discuss them here.

(VWO) takes six years and gives direct access to the academically oriented university, similar to universities in most of Europe.

While both universities and universities of applied science offer tertiary degrees, there are some important differences between them. First, universities of applied science have a strong vocational component. They prepare students for so-called 'professional' jobs, such as nurses or teachers at the primary or secondary level. They usually include mandatory internships at firms, often for periods up to half a year. University, on the other hand, is mostly academic and research-oriented. University is considered the highest education level in the Netherlands.[18] Second, a study at a university of applied sciences takes four years to complete, while regular university studies take four to six years, depending on the field of study. In both cases students immediately choose a field of study when they enroll. In principle each field is open to each student, although some technical studies require students to take additional courses in mathematics before they are allowed to enroll. Some fields of study (e.g. medicine) use a lottery because enrollment is larger than the number of available places. There is little overlap in the courses between the different fields, except for some common courses like basic statistics. This means that students graduating from different fields have acquired very different skill sets. Third, graduates from a university of applied science finish with a bachelor's degree, but most university students finish with a master's degree.[19] While it is possible for university students to enter the labor market after obtaining their bachelor's degree, this rarely happens. Close to 10% of university of applied science graduates continue to university to obtain a master's degree, usually after taking a bridge year to catch up with their academic skills. Around 90% of both university of applied science graduates and regular university graduates enter the labor market after finishing their degree.[20] Throughout the paper I will use the terms graduates from universities and academic graduates, and graduates from universities of applied science and vocational graduates interchangeably.

[18]The ISCED (2011) code for university bachelor degrees is 64, while the code for university of applied science bachelor degree is 65. A university master's degree corresponds to ISCED codes 74. A master's degree at a university of applied science has ISCED code 75 (Statistics Netherlands, 2011).

[19]Universities of applied science have started to offer master degrees in recent years, but these have low enrollment rates and are generally aimed at people who already have some work experience. Only 1% of graduates from universities of applied sciences in my sample have a master's degree.

[20]These numbers are based on public data from Statistics Netherlands.

3.2.4 Data sources and sample

I use administrative data from Statistics Netherlands on enrollment and graduation for all graduates in higher education from 1996 to 2012. The data contain detailed information on the type of programme followed - field of study and level - and the exact date of enrollment and graduation. These data can be merged at the individual level to other datasets using a coded social security number.[21] I merge administrative data on labor market status from 1999 to 2016 obtained from tax filings of employers. These contain the yearly gross wage and the number of days worked, which allows me to calculate the gross daily wage, my main dependent variable. The data also contain information on sector and an employer identifier. I take the main employer in a year to be the one where a worker earns the most in a year. I use these to identify when workers switch employer and sector, and to generate firm-level variables (see section 3.2.6). I obtain demographic characteristics by merging my data with municipal registries (GBA), which are available from 1995 onwards. These include age and gender. I also add information on social security claims and whether graduates work as self-employed. I do not have information on the level of the social security claims or the income earned as self-employed.

To obtain a sample of typical students, I restrict my sample in the following ways. First, I exclude students who first obtained a vocational (MBO) degree or a foreign degree before starting their higher education career. Second, I only include bachelor's, master's and equivalent degrees.[22] Third, I exclude everyone who graduated before the age of 20 or after the age of 30. Fourth, I exclude everyone who took shorter than three years or longer than seven years to obtain their degree. Fifth, I assume that students have entered the labor market if they have not been enrolled for at least 400 days after their graduation.[23] Finally, I drop workers who at some point earn less than 20% of the minimum daily wage, earn more than 700 euros a day

[21]The data are available via a secure connection to Statistics Netherlands for researchers who sign a confidentiality agreement.

[22]PhD's are not observed in the data. The only postgraduate degrees that are observed and dropped are professional postmaster degrees. These are usually only taken by people with some work experience. Postgraduate degrees are included when I examine whether economic conditions at graduation induce workers to obtain a higher degree. These results are reported in the Appendix.

[23]This assumption is necessary to define labor market entry. If graduates are not working and not enrolled for an additional degree, I do not observe them in the data. Hence, I assume that they are looking for a job during this period if it lasts for at least a year. The results do not depend on the exact length of this period.

(200,000 a year) or who stop filing taxes at some point because they have moved abroad.[24]

3.2.5 Constructing the change in field-specific employment

To construct the change in field-specific employment, I first calculate the share of workers with a degree from a given field of study and level (university or university of applied science) in each sector of industry. To this end, I use (the same) administrative data from tax filings that cover all workers from 1999 to 2016. I select all workers for whom I observe their highest degree and only keep those workers with a degree from higher education who are between 30 and 70 years old.[25] The remaining sample contains 11,934,327 worker-year observations. With 18 years of data, this is an average of 663,018 observations per year. Then I calculate for each year the share of workers with field of study f in sector of industry s. I use 33 fields of study (17 for universities of applied science and 16 for universities) and 44 sectors of industry. Fields of study are defined at the 2-digit ISCED97 level, where I combine some fields of study with a small number of observations.[26] Sectors of industry are defined using the 2-digit NACE Rev. 2. I join sectors where Statistics Netherlands does not provide official employment statistics or with a small number of graduates.[27] I then average the shares over all years to arrive at a time-invariant measure of the importance of each industry for each field. The change in employment in each industry is calculated using official statistics from the National Accounts from Statistics Netherlands covering 1995 - 2016.

Figure 3.1 shows the share of workers in the 10 largest fields of study (5 for universities of applied science and 5 for universities) working in different aggregated sectors. For many fields of study, public sectors such as government, health and education are important areas of employment. For others the private sector is more important. The figure highlights the variation in my employment decline measure that stems from time-invariant differences in employment patterns between fields of

[24] With these restrictions the sample shrinks from 1,000,929 graduates to 515,000 graduates. The bulk of the selection is due to dropping students with a vocational degree (283,784) and the subsequent selection on age and study duration (120,083).

[25] To prevent simultaneity bias, I drop workers younger than 30. The results are similar if I include all workers.

[26] For universities of applied sciences I join 42-46, 52 & 54, 62 & 64 and 85 & 86. For universities I combine 31 & 32, 52 & 54, 62 & 64 and 81 & 84-86. See Table A11 in the Appendix for all fields of study I use and the corresponding ISCED codes.

[27] I combine sectors 1-3, 6-9, 10-12, 13-15, 16-18, 19-23, 24-25, 26-27, 29-30, 31-32, 35-39, 50-51, 55-56, 58-60, 62-63, 69-71, 73-77, 80-82, 87-88, 90-92, 94-98.

study. The other part of the variation in the employment change measure derives from differences in employment changes between sectors. Figure 3.2 shows the variation in employment change for 9 aggregated sectors. The variation is largest for information and communication and business services. It is smallest for the public sector, which includes health and government.

Figure 3.1: The share of workers in each aggregated sector for the 5 largest fields of study at each level.

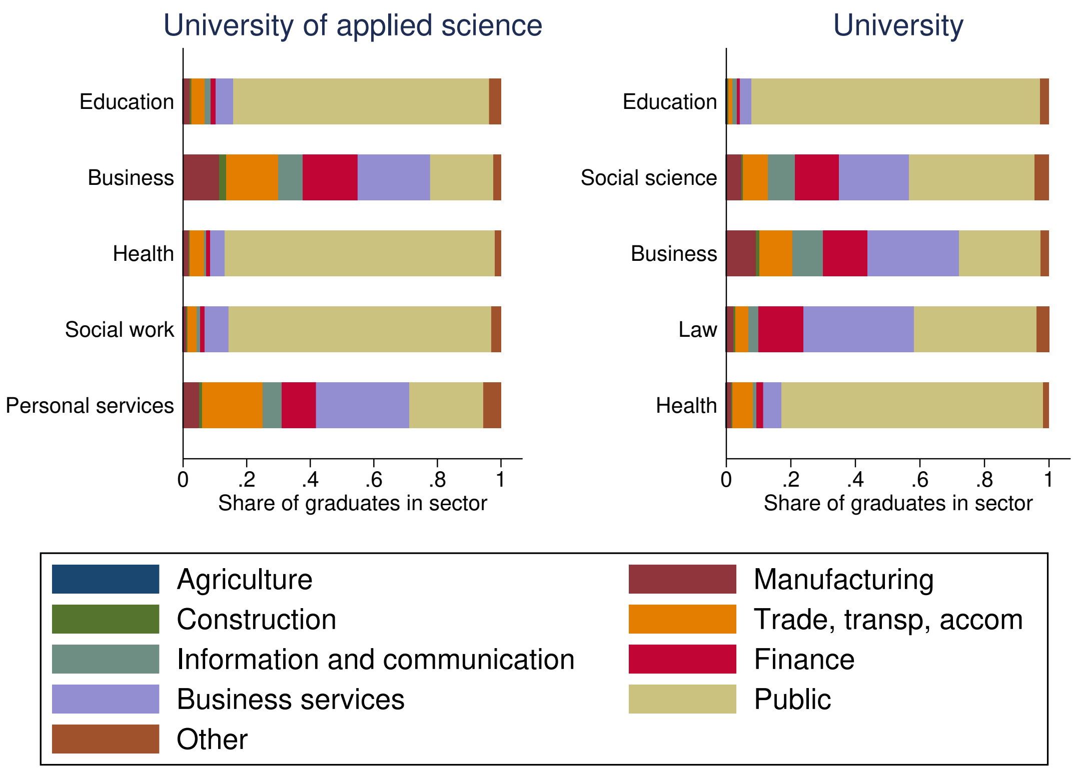

Notes: Sectors are aggregated at the NACE Rev2 main group level. Real estate (L) is combined with Business services (M-N) because it employs few high educated.
Source: Own calculations based on registration data from Statistics Netherlands.

3.2.6 Constructing other dependent variables

To construct firm-level measures I use the universe of Dutch firms from administrative linked employer-employee data for 1999 - 2016. For each year I select all workers who work at a firm on the 1st of October.[28] To ensure that I get an accurate picture of

[28]This is an arbitrary date, the results remain the same if I pick another date

Figure 3.2: Percentage change in employment for aggregated sectors.

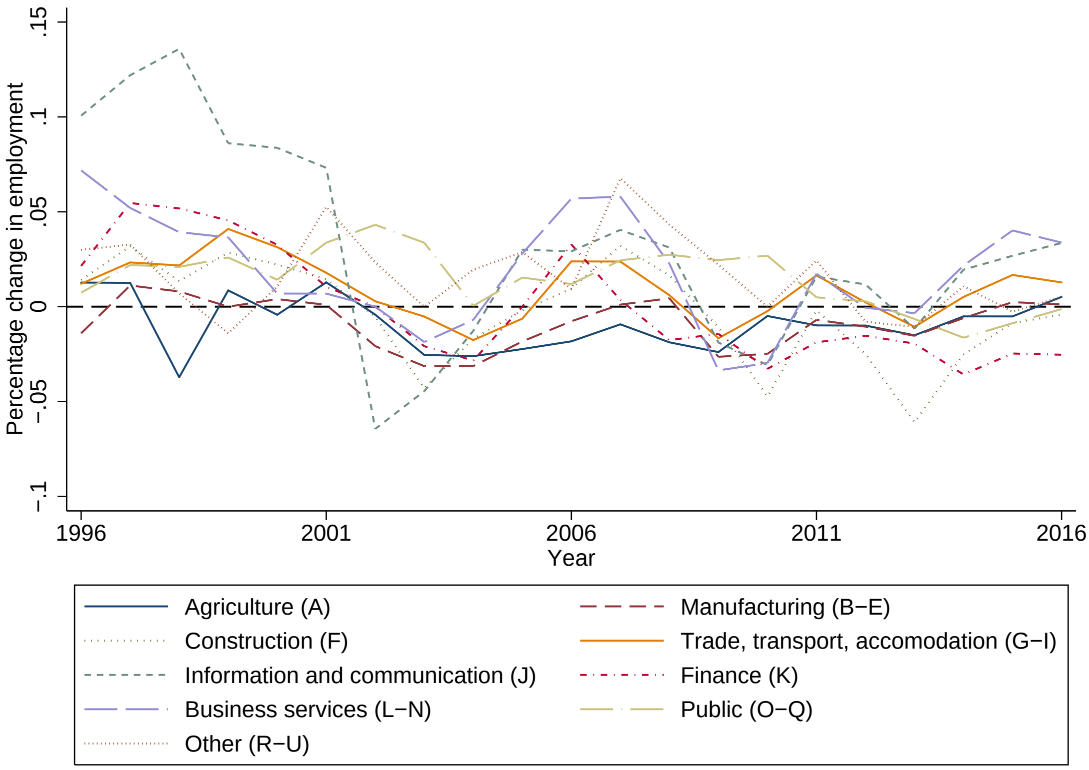

Notes: Sectors are aggregated at the NACE Rev2 main group level. Real estate (L) is combined with Business services (M-N) because it employs few high educated.
Source: Own calculations based on registration data from Statistics Netherlands.

the wages paid by a firm, I drop all workers who worked for less than 90 days at the firm in a year and all workers who earn less than 20% of the daily minimum wage or who earn more than 700 euros a day. I drop workers younger than 15 or older than 75. I also add the same 2-digit industry codes as used before. I calculate the average real yearly wage at the firm level. This is my main measure of firm quality. Measures of firm wage rank are constructed as in Haltiwanger et al. (2017). Within each 2-digit sector and year I calculate employment-weighted quintiles of the average gross real wage at the firm-level. I then classify high-rank firms as belonging to quintiles 4 and 5, medium-rank firms as belonging to quintiles 2 and 3 and low-rank firms as belonging to the first quintile.

To construct a measure of match quality, I use the same data as for constructing the field-specific employment. To prevent reverse causation, I keep all workers between 30 and 70 years old.[29] I then calculate the probability for a worker with a given degree f to work in a specific 2-digit sector s. Simply using these shares as an indicator for a match would lead to large sectors being overrepresented as good matches. Therefore, I normalize the measure by dividing these shares by the average probability for anyone with a higher education degree to work in sector s:[30]

$$M_f^s = \frac{S_f^s}{\sum_f S_f^s} \tag{3.3}$$

where M is the match quality of field of study f in sector of industry s and S is the share of workers in s with f. The match quality effectively gives the increase in probability to work in a s for workers with f compared to all higher educated workers. I then take the top 5 of sectors for each field of study as sectors where workers have a good match.[31]

Finally, switching firm or sectors for workers is defined as having a different firms or sector in year t compared to year $t-1$. All firms are contained within sectors, so workers can only switch sectors if they also switch firm.

3.2.7 Descriptive statistics

Table 3.1 gives descriptive statistics for my outcome and control variables. I work with an unbalanced panel of 3.7 million observations of more than 0.5 million unique individuals. With 18 cohorts, this amounts to an average of 30,365 individuals per cohort. 42% of the sample consists of university graduates, and 58% are graduates from universities of applied science. The average age at graduation is around 23 years for graduates from universities of applied science and almost 25 years for graduates from universities. Close to 60% of the sample is female. Since I have a sample of young workers, job mobility is high. 27% of the person-year observations involve a person switching firms and around 15% also switch sectors.[32] On average, close to half of vocational graduates and 42% of university graduates are in a good match. Most graduates are employed. Less than 2% are on benefits, while around 4% are self-employed.

[29] I also calculated the same measure using only younger workers between 20 and 35 years. The results are similar.

[30] This is similar to the share indicator in Liu et al. (2016).

[31] In Table A7 in the Appendix I present sensitivity analyses using other cutoffs. The results are very similar.

[32] Descriptives on mobility and firm rank for each year since graduation are reported in Tables A12 and A13, respectively.

There is substantial variation in economic conditions for different fields of study. Figure 3.3 reports the change in employment over time for the 5 largest fields of study for both universities of applied science and university. While the cyclical pattern can be observed for all fields, it is especially pronounced for business, personal services and law. These are fields that rely for a substantial part on the private sector for employment. Cyclical variation is much weaker for health and education, which rely more strongly on the public sector. Table A11 in the Appendix shows descriptives on the change in employment for each field of study. The strongest variation is found for technical studies, such as computing and engineering and manufacturing. This is probably related to the dot-com crisis in the early 2000s, as well as the strong employment growth of these fields. The weakest variation is observed for education and health.

Figure 3.4 show the wage-experience profiles for the cohorts in my sample. Starting wages (the dotted black line) differ quite strongly in line with the business cycle. However, wages seem to converge in the long run.

Figure 3.3: The change in employment over time for the 5 largest fields of study at each level.

(a) University of applied science (b) University

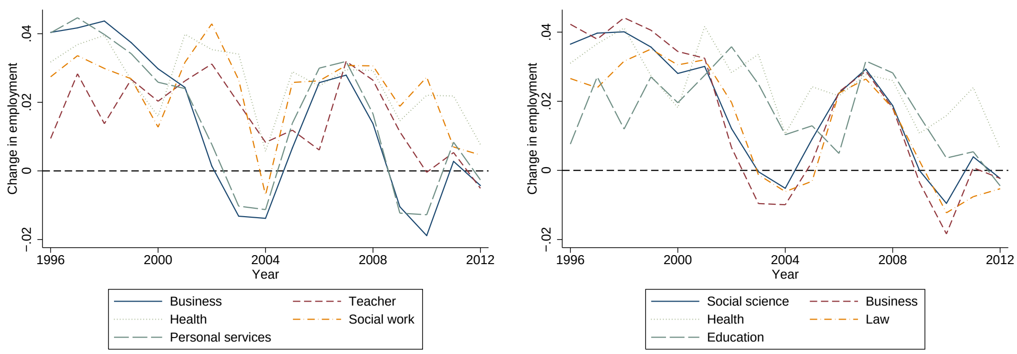

Source: Own calculations based on registration data from Statistics Netherlands.

3.3 Wage and employment effects of graduating during a recession

First I consider the effects of the change in employment at graduation on the log of daily wage. Since I am interested in the effect of graduating during a recession,

Figure 3.4: Wage-experience profiles for graduates from universities of applied science and university graduates.

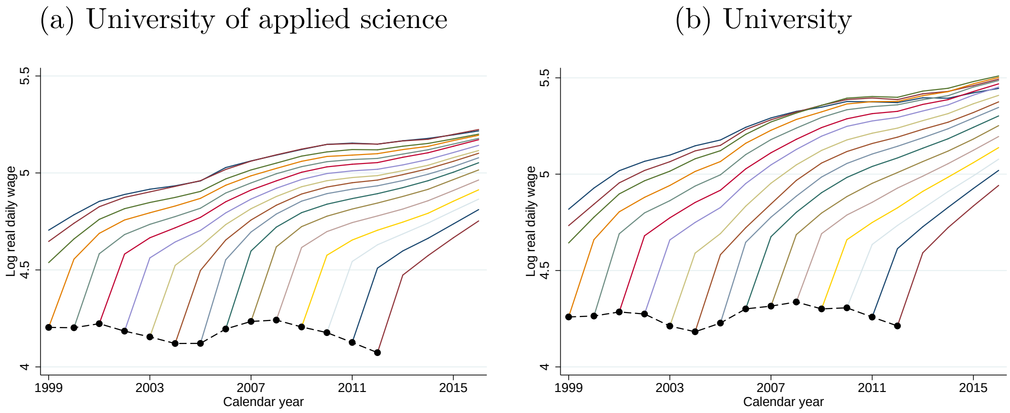

Source: Own calculations based on registration data from Statistics Netherlands.

to ease the interpretation of the results I will use the decline in employment as the main independent variable rather than the growth. Figure 3.5 plots the estimated coefficients from the IV specification with the dependent variable specified at the top.[33] Table A8 in the Appendix reports the estimates. The figures report the effect of a one percentage point decline in employment (around half of a standard deviation) on the specified outcome variable for the first 8 years after graduation. All models are separtely estimated for graduates from universities of applied science and graduates from university. Colored dots are significant at 5%, while white dots are not. Figure 3.5a shows that, conditional on employment, both vocational and academic graduates suffer substantial wage losses right after graduation if they started in a recession compared to their peers. For academic graduates the losses are close to 10% in the first year for each percentage point decline in field-specific employment, while the effect is around 6% for vocational graduates. Academic graduates then slowly catch up to their peers. While the coefficient remains negative up to 8 years after graduation, the effects are not statistically significantly different from zero after 6 years of experience. For vocational graduates the catch-up process is mostly concentrated in the second year after graduation, where the loss declines by around 50%. After the second year catching up slows down, but the estimated effect gradually declines to about -1% for each percentage point decline in employment. The estimated effects for vocational graduates remain significant up to 8 years

[33]The first stage estimates are repoted in Table A3 in the Appendix. The estimates are highly significant. They show that for both vocational and academic graduates the change in employment at graduation is strongly related to the change in employment in the nominal year of graduation.

Table 3.1: Descriptive statistics.

	University of applied science (1)	University (2)
Main independent variables		
Change in employment in year of graduation	1.7003	1.6153
	(1.8609)	(1.7235)
Change in employment in nominal year of graduation	1.7955	1.9403
	(1.8863)	(1.7473)
Main outcome variables		
ln(daily wage)	4.6982	4.8326
	(0.3975)	(0.4627)
Employed	0.9487	0.9410
	(0.2206)	(0.2356)
Self-employed	0.0402	0.0382
	(0.1965)	(0.1916)
On benefits	0.0199	0.0184
	(0.1396)	(0.1344)
Switch firm	0.2676	0.2891
	(0.4427)	(0.4533)
Switch 1-digit sector	0.1436	0.1601
	(0.3507)	(0.3667)
Switch 2-digit sector	0.1597	0.1743
	(0.3663)	(0.3794)
ln(average firm wage)	10.2659	10.4019
	(0.4031)	(0.4170)
Good match	0.4891	0.4203
	(0.4999)	(0.4936)
Control variables		
Female	0.5996	0.5691
	(0.4900)	(0.4952)
Age at graduation	23.2815	24.6928
	(1.7123)	(1.6922)
Observations	2,179,304	1,559,875
Number of individuals	298,946	216,054

Source: Own calculations based on registration data from Statistics Netherlands.

after graduation. The differences between vocational and academic graduates are significant for the first four years. However, for later years I can not reject the null hypothesis that the effects are equal.

These findings are generally robust to specification changes. Table A4 reports estimates using the national change in employment. The national estimates are smaller for both vocational and academic graduates, indicating that graduates respond much more strongly to conditions in labor markets where their studies prepared them for than to overall conditions. This could also explain why these estimates for academic graduates, although large, are not statistically significant. The same table also reports estimates using the sector-specific change in added value rather than employment as an input to my measure for economic conditions. The estimates are smaller than when using employment. This is expected, because there is generally some time lag between a loss in output and a reduction in employment. Nevertheless, the estimates are qualitatively in line with the main results. Table A5 reports OLS estimates. The estimates are smaller than the IV estimates, but the pattern is very comparable.[34]

One might also wonder whether the estimated effects are indeed linear as the specification used in this paper and in the literature assumes. Figure A1 in the Appendix reports marginal effects of a quadratic specification of the change in employment at graduation. The initial effects are larger for both academic and vocational graduates, but particularly for academic graduates. Nevertheless, the overall pattern remains the same. In fact, the catch-up seems to happen somewhat quicker for both groups when using this specification. Figure A2 compares the effects for graduates going into a downturn and graduates going into an upturn. A downturn is defined as a year-on-year increase in the (national) unemployment rate. This shows that the effects of a decline in employment also occur when going into an upturn, but that they are much stronger when going into a downturn. In particular again academic graduates suffer much stronger initial effects and their catch-up is slower than average. Finally, I also consider whether the effects differ across the two recessions - the Great Recession starting in 2008 and the dot-com crisis in the early 2000s - in my sample. These estimates show that, again, the initial effects are much stronger for those graduating in the stronger recession from 2008 onwards. Note that the estimates for the later years since graduation should be interpreted with caution, because I can only follow the 2007 and 2008 cohorts for the full 8 years.

Labor market conditions at graduation might also affect labor supply decisions. And if lower-skilled workers are more likely to quit looking for a job in a recession, I might underestimate the effect of the unemployment rate at graduation on wages. It is therefore instructive to also consider effects on employment and other labor market outcomes.

Figure 3.5b reports results for employment. For both vocational and academic graduates I find small declines in employment probabilities right after graduation.

[34]Table A6 reports further robustness checks. First, it presents estimates from the main specification but restricted to students who did not switch from their initial track assigned to them at the end of primary school. The results are very similar to the main results. This alleviates some concerns about students selecting themselves in the academic or vocational track. Second, it presents estimates only including workers who report positive earnings in each year since graduation. The estimates are again very similar, indicating that there is no selective dropping out of the labor market due to bad starting conditions.

Figure 3.5: Estimated effects of the decline in field-specific employment on wage and employment status.

(a) ln(daily wage)

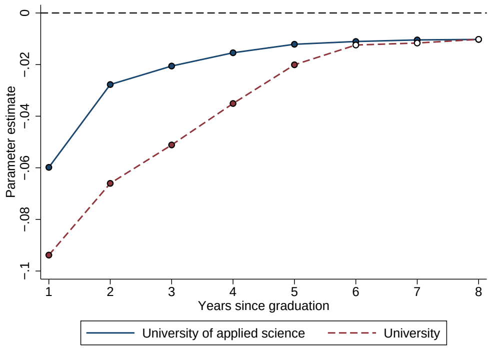

(b) Employment

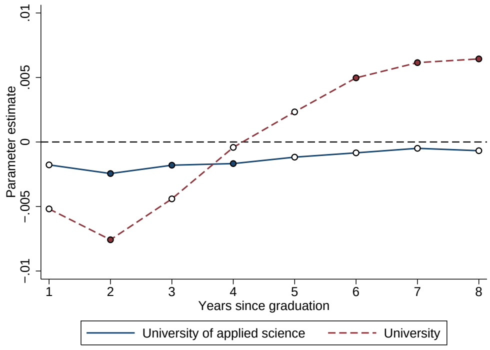

(c) Self-employment

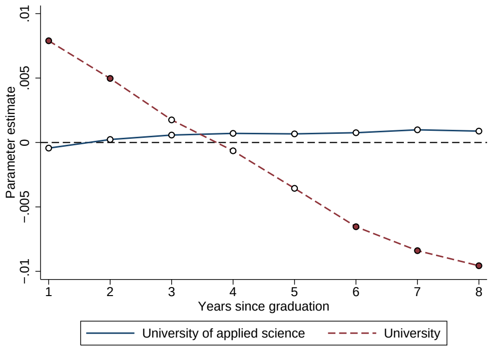

(d) On benefits

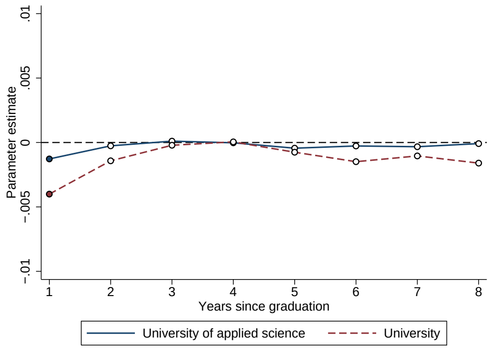

Notes: The figure reports estimates of the effect of the decline in field-specific employment on ln(daily wage), a dummy for employment, a dummy for self-employment and a dummy for being on benefits at different years since graduation. Coefficients are obtained from IV estimates of equation 3.2 where the decline in field-specific employment at graduation is instrumented with the decline in field-specific employment at the nominal moment of graduation. Colored dots are statistically significant at at least 5% and white dots are not. Parameter estimates are reported in Table A8 in the Appendix.
Source: Own calculations based on registration data from Statistics Netherlands.

After about 3 to 4 years on the labor market however, the effects disappear. For each percentage point decline in employment, the probability to be employed declines by about 0.5% for academic graduates and by 0.2% for vocational graduates. The effects do not differ significantly however. For university graduates I find small positive effects in later years.

If it is difficult to find a job, graduates might substitute regular employment with self-employment. In addition, for some fields of study, such as arts, self-employment is very common. On the other hand, graduates might be more inclined to take the risk of becoming self-employed in good economic conditions (Beiler, 2017). Figure 3.5c shows the estimates on the probability to be self-employed. I find that the probability to be self-employed increases in the first few years after graduation in bad economic conditions for university graduates. These results are in contrast to Beiler (2017), who finds for Germany that the probability to be self-employed declines for recent graduates in bad economic conditions. This suggests that in the Netherlands self-employment acts as a substitute for regular employment when it is more difficult to find a job. I find negative effects for university graduates in later years. Together with the estimates for employment, this suggests that they switch from employment to self-employment. For vocational graduates I find no effect of economic conditions on self-employment.

Finally, I examine whether graduates are more likely to be on benefits. This includes unemployment benefits, welfare and disability benefits. Figure 3.5d reports the results. I find only an effect for the first year after graduation, and it is a small negative estimate. Overall, graduates do not seem to rely on benefits to supplement lost labor income.

In all, the results suggest that vocational graduates have an easier transition to the labor market in bad economic conditions than academic graduates. Both their wage and employment losses compared to their peers are smaller than for academic graduates. This confirms the hypothesis that vocational education eases the transition to the labor market in a recession. Academic graduates ultimately catch up to those who graduated in good times in about five years. For vocational graduates, on the other hand, I find persistent wage losses at about -1% for at least the first 8 years of their career. I also find evidence that the effects of graduating in a recession are not linear. In particular for academic graduates I find that the initial effects are much stronger in a downturn and for those who graduated in the Great Recession compared to earlier cohorts. I will now turn to the mechanisms driving the recovery process.

3.4 Mechanisms of recovery

Young workers can recover from initial wage losses due to bad starting conditions both within firms and between firms. Within firms human capital theory helps to explain the recovery process. Building firm and sector-specific human capital increases productivity and will lead to wage increases. If learning is concave, workers who start in worse positions will eventually catch up with those who started in better positions. An alternative model to explain both the initial losses and the subsequent recovery is long-term wage contracting. In this model workers and firms agree on a contract wage when an employment relationship starts. If labor market conditions are bad during the start, the contracted wage will be lower. If conditions improve and workers are mobile, the firm has to renegotiate the wage to keep the worker from moving (Harris and Holmstrom, 1982; Beaudry and DiNardo, 1991). Both models imply that workers who start in a recession will primarily improve their wage within the firm, not by moving to another firm. The human capital view explains this through building firm-specific human capital.[35] The long-term contracting view relies on renegotiated contract wages.

Catch up between firms occurs through job mobility. It is a costly process of finding the right job. This can be explained by search theory, in particular by models that feature on the job search. As already highlighted by Topel and Ward (1992), frequent job switching is an integral part of most early careers. Through job shopping, young workers search for a good match and they experience wage gains. Indeed, wage gains in early careers are for a substantial part explained by labor market frictions (Topel and Ward, 1992; Van der Klaauw and van Vuuren, 2010).

When graduates enter the labor market in a recession, there are fewer vacancies. Furthermore, there is substantial evidence that the quality of vacancies is lower (McLaughlin and Bils, 2001; Martins et al., 2012; Moscarini and Postel-Vinay, 2016; Haltiwanger et al., 2017). This suggests that workers are less likely to find a good match when they start in a recession. This is confirmed by recent evidence that young workers who start in a recession are more likely to be mismatched (Liu et al., 2016; Altonji et al., 2016).

A useful model to think about this process is a job ladder model. The idea of this type of model is that firms are ranked in terms of the wages they pay. High-paying firms, who pay more because they are more productive, are high up the job ladder, whereas low-paying firms who are less productive are low-ranked. Workers search for a job both while unemployed and on the job. Workers only accept an offer if it is better than their current offer, and hence, through on the job search, workers move up the job ladder to high-paying firms. More productive firms grow by poaching workers from low-paying firms, while low-paying firms more often hire from unemployment. During recessions, the process of moving up the job ladder slows

[35] Of course, if the initial sector or firm is a particularly bad match, it will be helpful to switch to a better matched sector or firm. Otherwise, workers might be stuck in a bad match because they start developing specific human capital for the tasks in that match (Gibbons and Waldman, 2004, 2006). Nevertheless, the theory predicts that workers will only start catching up once they start developing their human capital.

down (Moscarini and Postel-Vinay, 2008, 2012, 2013, 2016). High-paying firms, who were less restricted in hiring during the previous boom because they could poach from low-paying firms, now have more employment to shed than low-paying firms. Furthermore, due to high unemployment in a recession, all firms find it easier to fill vacancies. This primarily affects low-paying firms however, since they more often hire from unemployment. Low-paying firms will therefore hire more during recessions than high-paying firms. This also means that the probability of moving up the job ladder declines (Haltiwanger et al., 2017).

Recovery through productive on the job search implies that workers who start during a recession will be more likely to switch firms and sectors until they catch up. In addition, the dynamic job ladder model implies that workers are more likely to start at low-paying firms. Finally, if primarily high-paying firms sharply reduce hiring, the probability of an initial mismatch for high educated workers could also increase.[36] I will examine these mechanisms in this section.

3.4.1 Firm and sector mobility

I first consider the quality of firms workers who graduated in a recession start at. I use the log real mean yearly wage of a firm as the main indicator for firm quality. This is a simple measure of "firm quality". Higher paying firms are thought to be of higher quality, and could be a better match, especially for high-educated graduates.[37] Figure 3.6a shows that both vocational and academic graduates tend to start at lower-paying firms if they graduated in a recession. For academic graduates I find that for every percentage point decline in employment at graduation, the average wage paid by the firm is 5% lower compared to their peers. The effects are somewhat smaller for vocational graduates at about 3.5%. However, these effects are not significantly different from each other. Both academic and vocational graduates recover, until there are no significant differences anymore at the fourth or fifth year after graduation.

The steady increase in firm quality already suggests that job mobility increases. Figures 3.6b to 3.6d confirm this. They show estimates of the effect of the decline in employment at graduation on the probability to switch firm, 1-digit and 2-digit sector. The estimates show that for both graduates from universities and universities

[36] This is not necessarily implied by the model since it does not feature worker and job heterogeneity.

[37] Results are very similar if instead of the mean wage I use a wage measure adjusted for firm-level characteristics.

of applied science job mobility increases in the first 5 years if they started in a bad labor market. The estimates suggest that the effects are somewhat stronger for graduates from university, who also suffered stronger wage losses, but they are not significantly different from each other. After about five years on the labor market, firm and sector mobility is similar to those who graduated in a boom. This is also the point when most of the wage losses have been recoverd. For graduates from universities of applied science the estimates for sector mobility remain significant up to 7 years after graduation.[38]

These findings are in line with earlier findings on the more flexible labor market in Canada (Oreopoulos et al., 2012). Job mobility, and hence on the job search, plays an important role in recovery from a bad starting position for both academic and vocational graduates.[39]

3.4.2 Match quality and the job ladder

I will now consider the effects of the decline in employment at graduation on match quality and firm rank (as defined in section 3.2.6). Figure 3.7 reports the estimation results. Figure 3.7a presents the effect on match quality. I find that for both academic and vocational graduates the match quality is lower if they started in a recession. For academic graduates match quality quickly improves, and the effects are no longer significant after year 3. For vocational graduates, however, match quality remains consistently lower than for those who started in good times.[40] Figures 3.7b to 3.7d show the effect of the decline in employment at graduation on the probability of

[38]Some of the estimates for switching sector are larger than for switching firms, while firms are contained within a sector. The reason is that these estimates present the effect on the probability of switching relative to workers with the same degree who started in better times. They appear just as likely to switch firms as those who started in bad times after a few years on the labor market, but are less likely to switch sectors.

[39]Earlier literature on job mobility for vocational and general educated workers focuses mostly on graduates from apprenticeships in German-speaking countries or secondary vocational education (e.g. Acemoglu and Pischke, 1998; Korpi and Mertens, 2003; Von Wachter and Bender, 2006; Göggel and Zwick, 2012; Dustmann and Schönberg, 2012; Fitzenberger et al., 2015). This literature often finds negative wage returns to job switching after finishing apprenticeships, although some studies find no effects or small positive effects. However, for apprenticeships the losses are likely driven by the loss of firm-specific human capital. Korpi and Mertens (2003) show that the probability to switch jobs is lower for those with more general secondary education than for those from an apprenticeship, but they find no differences in firm and industry mobility. There is as far as I'm aware no evidence on differences in job mobility for tertiary vocational and academic graduates.

[40]Table A7 reports robustness checks using other cutoffs for the match quality indicator or other match quality measures. The results are generally very robust. I find that in each case vocational graduates are more likely to remain mismatched, while for academic graduates the mismatch is recoverd after 3 to 5 years on the labor market.

Figure 3.6: Estimated effects of the decline in field-specific employment on mean firm wage and job mobility.

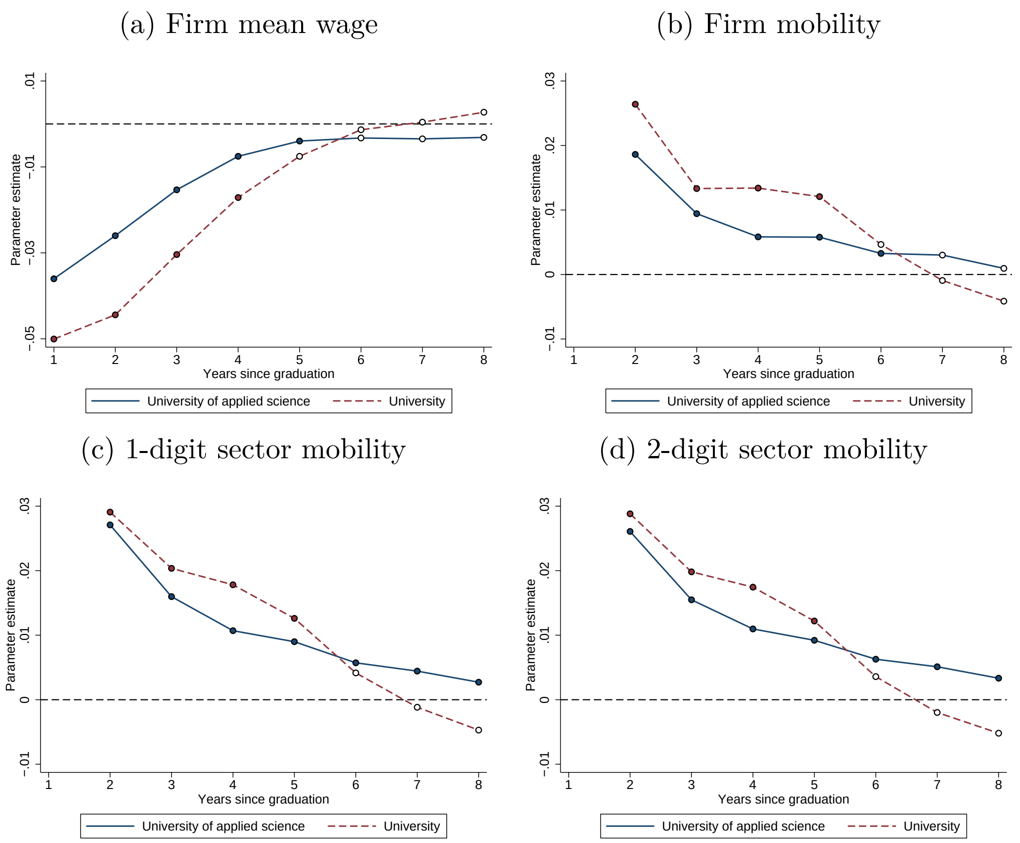

Notes: The figure reports estimates of the effect of the decline in field-specific employment on ln(mean firm wage), a dummy for switching firm, a dummy for 1-digit sector switch and a dummy for switching 2-digit sector at different years since graduation. Coefficients are obtained from IV estimates of equation 3.2 where the decline in field-specific employment at graduation is instrumented with the decline in field-specific employment at the nominal moment of graduation. Colored dots are statistically significant at at least 5% and white dots are not. Parameter estimates are reported in Table A9 in the Appendix.
Source: Own calculations based on registration data from Statistics Netherlands.

working in a low-, medium- or high-ranked firm. Consistent with the previous results, I find that the probability to work in a low or medium ranked firm increases if workers started in bad conditions, while the probability to work in a high-ranked firm declines. This is in line with the dynamic job ladder model, which predicts that lower-ranked firms are more likely to hire in recessions. Through job mobility, graduates eventually recover and climb the job ladder. Nevertheless, they remain more likely to work in medium-ranked firms and vocational graduates are significantly more likely than their peers to get stuck in low-ranked firms. This could at least partly explain the persisitently significant wage losses for vocational graduates.

Figure 3.7: Estimated effects of the decline in field-specific employment on match quality and firm rank.

(a) Match quality

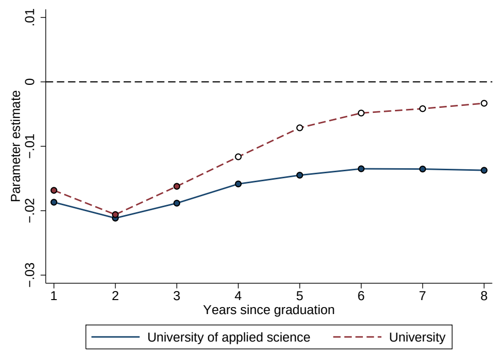

(b) Low-ranked firm

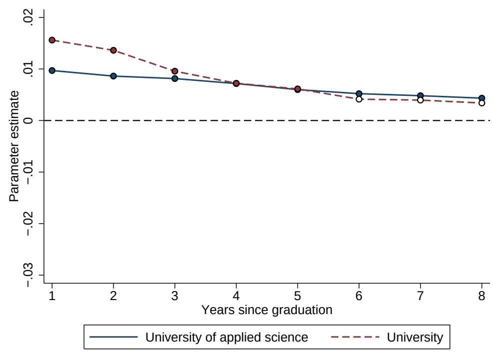

(c) Medium-ranked firm

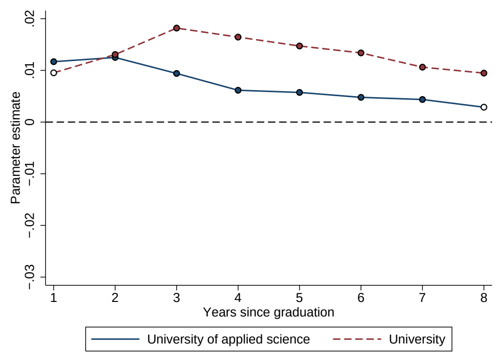

(d) High-ranked firm

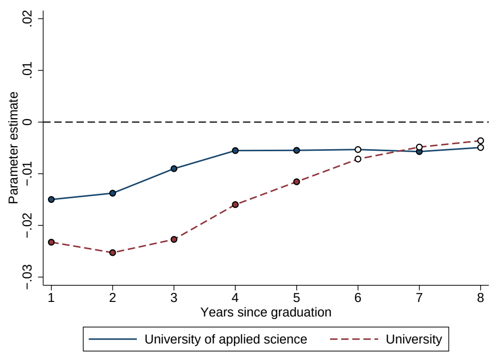

Notes: The figure reports estimates of the effect of the decline in field-specific employment on a dummy for being mismatched and dummies for working in a low-, medium- and high-ranked firm at different years since graduation. Coefficients are obtained from IV estimates of equation 3.2 where the decline in field-specific employment at graduation is instrumented with the decline in field-specific employment at the nominal moment of graduation. Colored dots are statistically significant at at least 5% and white dots are not. Parameter estimates are reported in Table A10 in the Appendix.
Source: Own calculations based on registration data from Statistics Netherlands.

3.4.3 Returns to job mobility

Up until this point the evidence on recovery via job mobility has been indirect. I will now consider some direct evidence by looking at the wage returns to job mobility. To estimate the returns, I augment my baseline model with a dummy variable for mover status ($Mover$) and interact it with the change in employment at graduation (e_cf). The dependent variable is the change in log daily wage ($\Delta ln(wage)_{it}$, defined as $ln(wage)_{it} - ln(wage)_{it-1}$):

$$\Delta ln(wage)_{it} = \alpha + \beta_{exp} exp_{it} * e_{cf} + \gamma_1 Mover + \gamma_2 Mover * e_{cf} + \eta_i + \tau_t + \varepsilon_{it}, \tag{3.4}$$

where $Mover$ is 1 if a worker switches firm (or sector) and 0 otherwise. Obviously, job mobility is endogenous. Workers who move are probably those who benefit from making a move. Therefore, these estimates are likely an upper bound of the actual effect of switching firms or sectors. To deal with the bias from endogenous job mobility as best as possible, I include individual fixed effects η_i to control for time-invariant individual factors (Von Wachter and Bender, 2006; Del Bono and Vuri, 2011). The other variables are defined as before.[41] Parameter γ_1 gives an estimate of the effect of moving to another job on changes in log daily wage and γ_2 gives the differential effect of changing jobs due to the change in employment at graduation. If those who graduate during a recession have higher returns to job mobility than those who graduate during a boom, I would expect a positive estimate for γ_2.

Table 3.2 reports the estimation results. First consider the estimates for vocational graduates in columns (1) to (3). Column (1) shows that if vocational graduates switch firm, they gain almost 2.5 log points in wage. Those who started in a recession gain around 0.8 log point per percentage point decline in employment at graduation. The gains when switching sector are higher at close to 4 log points (columns 2 and 3). These likely reflect the gain from switching to a better matched sector. The gains for those graduating in a recession are also larger at around 1 log point per percentage point declinie in employment at graduation. Now consider the results for academic graduates. Column (4) shows that academic graduates gain almost 5 log points when they switch firm. The gains from switching sector are also larger than for vocational graduates, at around 6.5 log points. The differences between vocational and academic graduates are highly significant. However, the gains for those who started in a recession are similar for academic and vocational graduates, at about 1 log point per percentage point decline in employment at graduation.

Taken together, these results suggest that job mobility plays a critical role in recovering from initial wage losses for both academic and vocational graduates who start in a recession. They are both more likely to switch firms and sectors, and when they do switch, they gain more than their counterparts who started in a boom. The results are in line with a model that emphasizes search on the job as a mechanism of wage gains. Consistent with the dynamic job ladder model, graduates are more likely to start in low-ranked firms in a recession and gradually move to higher ranked

[41] The cohort and field of study fixed effects drop out due to the individual fixed effects.

Table 3.2: Fixed effects estimates of the wage return to firm and sector mobility.

	University of applied science			University		
	Firm mobility	1-digit sector mobility	2-digit sector mobility	Firm mobility	1-digit sector mobility	2-digit sector mobility
	(1)	(2)	(3)	(4)	(5)	(6)
Switch	0.0247***	0.0393***	0.0382***	0.0489***	0.0655***	0.0661***
	(0.0024)	(0.0026)	(0.0025)	(0.0026)	(0.0033)	(0.0032)
Switch x emp decline	0.0077***	0.0097***	0.0097***	0.0107***	0.0097***	0.0098***
	(0.0011)	(0.0012)	(0.0012)	(0.0015)	(0.0017)	(0.0017)
N	1,757,362	1,757,362	1,757,362	1,242,155	1,242,155	1,242,155

Notes: The table reports the effect of switching firm or sector on the year-on-year change in log daily wage. Regressions include individual fixed effects and calendar year fixed effects. They are separately estimated for graduates from universities of applied science and university. Standard errors in parentheses are clustered at the level of graduation cohort and field of study. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
Source: Own calculations based on registration data from Statistics Netherlands.

firms. Interestingly, while switching sectors solves the initial mismatch for academic graduates, vocational graduates remain in sectors that are not typical for their field of study. This could at least partly explain the persistent wage losses for vocational graduates.

3.5 Conclusion

In this paper I examined the early career effects of graduating in a recession for vocational and academic graduates from Dutch higher education. I used cohorts from 1996 to 2012 and followed them on the labor market from 1999 to 2016. Exploiting field-specific differences in the change in employment at graduation, I find that academic graduates suffer an initial 10% lower wage per percentage point decline in employment at graduation. The effects gradually decline and fade out after about five years on the labor market. The initial wage losses for vocational graduates are significantly smaller at close to 6% for each percentage point decline in field-specific employment at graduation. However, they remain persistent at around -1% up until 8 years after graduation.

The main mechanism driving the initial losses is that graduates start working at employers who pay less and are lower on the job ladder. Through upwards job mobility they catch up to those who graduated in good times. Both groups of graduates are more likely to switch firms and sectors, and when they do switch, they gain more than their counterparts who started in a boom. Academic graduates gain significantly more when switching than vocational graduates. Both are also more likely to be mismatched in their early career. While for academic graduates switching sectors solves the initial mismatch, vocational graduates remain in sectors that are not typical for their field of study. This could at least partly explain the persistent wage losses for vocational graduates. To conclude, the transition from education to the labor market in a recession is easier for vocational graduates, but academic graduates ultimately catch up quicker to their peers who started in better times than vocational graduates.

A Appendix

A.1 Selection bias

The main analyses use an instrument to take into account that students might postpone graduation in a bad labor market. One way in which students might select themselves onto the labor market is in their choice of obtaining a higher or additional degree. Students who graduate during a recession might face lower opportunity costs of staying in school and thus are more likely to obtain an additional degree. Table A1 shows the estimated relation between the change in field-specific employment at graduation (the first level mentioned in each column) and the probability to obtain an additional degree (the second level mentioned). I estimate a simple linear probability model that relates a dummy variable indicating whether a student pursued an additional degree to the change in field-specific employment measured at graduation of the first level and the same set of control variables as included in the other specifications. Note that students with a degree from a university of applied science can pursue a bachelor's or master's degree at the their own level, or a bachelor's or master's degree at the university level. University students can pursue an (additional) master's degree at the university level. I find that both graduates from universities of applied science and graduates from universities are more likely to pursue a degree at the university level if employment in their field declined at graduation.

Table A2 provides further evidence of selection. It reports the effect of the decline in employment at graduation on the composition of the graduation cohorts. Columns (1) and (5) report the effect on the share of women. I find that cohorts from universities of applied science are less likely to be female when field-specific employment is higher. Columns (2) and (6) report the average age. I find no significant effect. Columns (3) and (7) report effects on the share of graduates 28 or older. This indeed increases for cohorts in universities of applied science. Finally, I find no effect on the share of workers 24 or younger.

Table A1: OLS estimates of the effect of the decline in employment at graduation on the probability of obtaining an additional degree.

	University of applied science to university of applied science	University of applied science to university	University to university
	(1)	(2)	(3)
Decline in employment at graduation	−0.0007	0.0052***	0.0020**
	(0.0005)	(0.0013)	(0.0009)
N	367,137	367,137	196,849

Notes: The table reports the effect of the decline in employment at graduation of the first degree indicated on the probability to obtain an additional degree at the second level indicated. The OLS regressions include field of study and graduation year fixed effects and controls for gender and age at graduation of the initial level. Standard errors in parentheses are clustered at the level of graduation cohort and field of study. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
Source: Own calculations based on registration data from Statistics Netherlands.

Table A2: OLS estimates of the effect of the change in employment at graduation on the composition of the graduation cohort.

	University of applied science				University			
	Share women	Average age	Share 28 or older	Share 24 or younger	Share women	Average age	Share 28 or older	Share 24 or younger
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Decline in employment at graduation	−0.0030*	−0.0017	0.0005**	0.0009	−0.0015	0.0127	0.0026	−0.0056
	(0.0018)	(0.0052)	(0.0003)	(0.0010)	(0.0020)	(0.0207)	(0.0019)	(0.0052)
N	298,946	298,946	298,946	298,946	216,054	216,054	216,054	216,054

Notes: The table reports the effect the percentage decline in employment at graduation the share of women, the average age and the share of 28 and older and 24 and younger in the graduation cohort. The OLS regressions include graduation year and field of study fixed effects and are separately estimated for cohorts from universities of applied sciences and universities. Standard errors in parentheses are clustered at the level of graduation cohort and field of study. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Source: Own calculations based on registration data from Statistics Netherlands.

Table A3: First stage estimates for IV estimates.

	University of applied science		University	
	First stage estimate	F-statistic	First stage estimate	F-statistic
	(1)	(2)	(3)	(4)
Effect at year of potential experience				
1	0.6699***	570.39	0.4179***	173.56
	(0.0285)		(0.0280)	
2	0.7051***	368.27	0.4786***	177.86
	(0.0284)		(0.0308)	
3	0.7281***	534.17	0.5256***	184.04
	(0.0266)		(0.0313)	
4	0.7566***	422.37	0.5559***	182.85
	(0.2780)		(0.0311)	
5	0.7579***	683.23	0.5571***	193.48
	(0.0278)		(0.0313)	
6	0.7525***	715.86	0.5399***	246.29
	(0.0293)		(0.0335)	
7	0.7639***	716.85	0.5332***	205.98
	(0.0298)		(0.0360)	
8	0.7388***	481.58	0.4975***	179.28
	(0.0334)		(0.0370)	
N	2,067,535		1,467,864	

Notes: The first stage regressions include fixed effects for potential experience, calendar year, field of study and graduation year. Demographic controls are age at graduation and gender. Standard errors in parentheses are clustered at the level of graduation cohort and field of study. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
Source: Own calculations based on registration data from Statistics Netherlands.

A.2 Other rubustness checks

Figure A1: Marginal effects of the decline in field-specific employment on ln(daily wage) using a quadratic specification for the decline in field-specific employment.

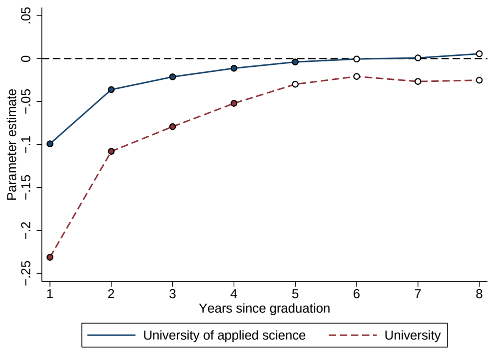

Notes: The figure reports marginal effects of the effect of the decline in field-specific employment on ln(daily wage) using a quadratic specification for the decline in field-specific employment. Marginal effects are obtained from IV estimates of equation 3.2 where the decline in field-specific employment at graduation is instrumented with the decline in field-specific employment at the nominal moment of graduation. Colored dots are statistically significant at at least 5% and white dots are not.
Source: Own calculations based on registration data from Statistics Netherlands.

Figure A2: Estimated effects of the decline in field-specific employment on ln(daily wage) for workers graduating in a downturn or an upturn and graduating in the Great Recession or before.

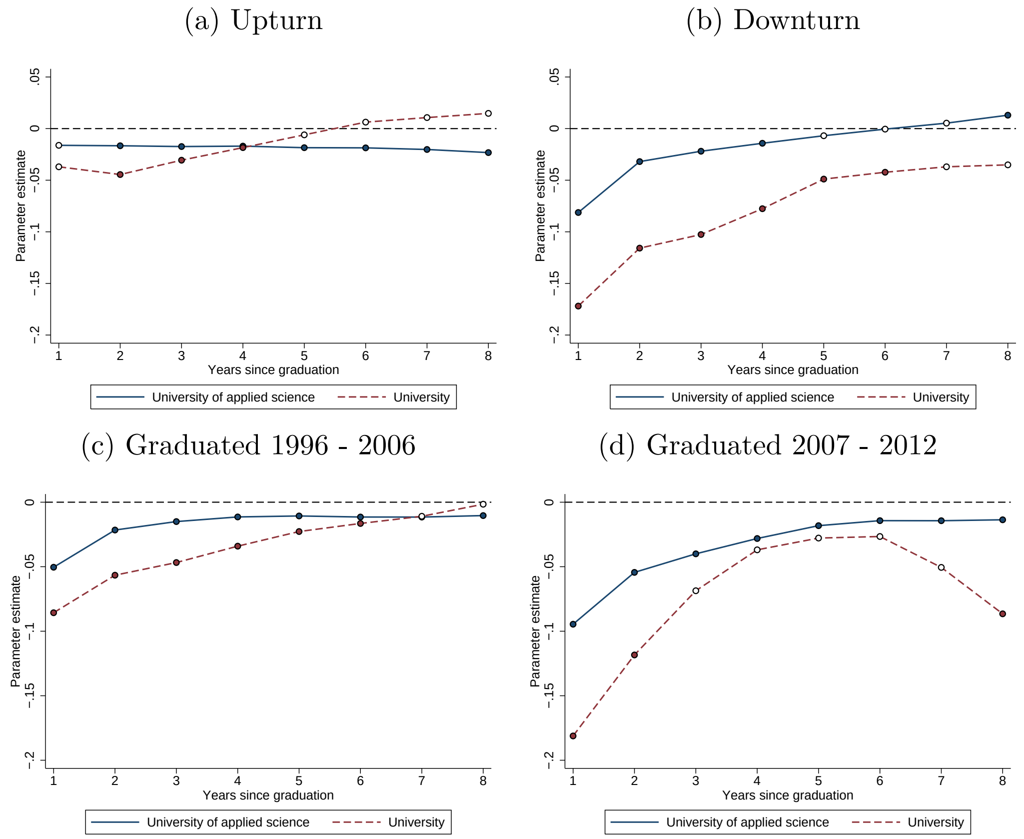

Notes: The figure reports estimates of the effect of the decline in field-specific employment on ln(daily wage). Coefficients are obtained from IV estimates of equation 3.2 where the decline in field-specific employment at graduation is instrumented with the decline in field-specific employment at the nominal moment of graduation. Colored dots are statistically significant at at least 5% and white dots are not.
Source: Own calculations based on registration data from Statistics Netherlands.

Table A4: IV estimates of the effect of two different indicators for economic conditions at graduation on ln(daily wage). The first measure uses the national change in employment. The second uses the sector-specific change in added value as input in calculating field-specific economic conditions.

	University of applied science		University	
	National change in employment (1)	Added value instead of employment (2)	National change in employment (3)	Added value instead of employment (4)
Effect at years since graduation				
1	-0.0344^{***}	-0.0214^{***}	-0.0421	-0.0309^{**}
	(0.0115)	(0.0068)	(0.0298)	(0.0136)
2	-0.0093^{***}	-0.0110^{***}	-0.0234	-0.0218^{***}
	(0.0021)	(0.0031)	(0.0147)	(0.0084)
3	-0.0034^{***}	-0.0110^{***}	-0.0185	-0.0186^{***}
	(0.0012)	(0.0023)	(0.0132)	(0.0071)
4	-0.0013	-0.0120^{***}	-0.0103	-0.0143^{**}
	(0.0017)	(0.0019)	(0.0086)	(0.0057)
5	-0.0024	-0.0141^{***}	-0.0004	-0.0098^{**}
	(0.0021)	(0.0017)	(0.0041)	(0.0047)
6	-0.0052^{***}	-0.0170^{***}	0.0047	-0.0086^{*}
	(0.0018)	(0.0020)	(0.0041)	(0.0044)
7	-0.0062^{***}	-0.0193^{***}	0.0043	-0.0102^{**}
	(0.0017)	(0.0023)	(0.0049)	(0.0042)
8	-0.0052^{*}	-0.0212^{***}	0.0045	-0.0112^{***}
	(0.0029)	(0.0027)	(0.0061)	(0.0043)
N	2,067,535	2,067,535	1,467,864	1,467,864

Notes: The table reports the effect of the percentage decline in employment (or added value) at graduation on ln(daily wage). Columns (1) and (3) use the national change in employment rather than the field-specific change. Columns (2) and (4) use the sector-specific change in added value (in real 2010 prices) as input for calculating the economic conditions measure rather than the change in employment. Coefficients obtained from IV regressions where the percentage decline in employment (or added value) at graduation is instrumented with the percentage decline in employment (or added value) in the nominal year of graduation. Regressions include fixed effects for potential experience, calendar year, field of study and graduation year. Demographic controls are age at graduation and gender. They are separately estimated for graduates from universities of applied science and university. Standard errors in parentheses are clustered at the level of graduation cohort and field of study. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
Source: Own calculations based on registration data from Statistics Netherlands.

Table A5: OLS estimates of the effect of the field-specific decline in employment on ln(daily wage).

	University of applied science (1)	University (2)
Effect at years since graduation		
1	-0.0388^{***}	-0.0378^{***}
	(0.0054)	(0.0067)
2	-0.0207^{***}	-0.0299^{***}
	(0.0028)	(0.0037)
3	-0.0159^{***}	-0.0223^{***}
	(0.0022)	(0.0031)
4	-0.0120^{***}	-0.0136^{***}
	(0.0018)	(0.0026)
5	-0.0095^{***}	-0.0061^{**}
	(0.0019)	(0.0026)
6	-0.0084^{***}	-0.0031
	(0.0023)	(0.0029)
7	-0.0078^{***}	-0.0025
	(0.0027)	(0.0033)
8	-0.0081^{**}	-0.0020
	(0.0033)	(0.0036)
N	2,067,535	1,467,864

Notes: The table reports OLS estimates of the effect of the percentage decline in employment at graduation on ln(daily wage). Regressions include fixed effects for potential experience, calendar year, field of study and graduation year. Demographic controls are age at graduation and gender. They are separately estimated for graduates from universities of applied science and university. Standard errors in parentheses are clustered at the level of graduation cohort and field of study. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
Source: Own calculations based on registration data from Statistics Netherlands.

Table A6: IV estimates of the effect of the decline in employment at graduation on ln(daily wage) for those who remained within their initial track from secondary education and for those who are observed as employed for each year since graduation.

	Never switched track from secondary school		Always employed	
	University of applied science	University	University of applied science	University
	(1)	(2)	(3)	(4)
Effect at years since graduation				
1	-0.0629^{***}	-0.0826^{***}	-0.0586^{***}	-0.0875^{***}
	(0.0083)	(0.0164)	(0.0073)	(0.0174)
2	-0.0282^{***}	-0.0543^{***}	-0.0276^{***}	-0.0622^{***}
	(0.0043)	(0.0110)	(0.0040)	(0.0132)
3	-0.0203^{***}	-0.0365^{***}	-0.0200^{***}	-0.0482^{***}
	(0.0034)	(0.0093)	(0.0030)	(0.0123)
4	-0.0145^{***}	-0.0229^{***}	-0.0147^{***}	-0.0329^{***}
	(0.0028)	(0.0079)	(0.0024)	(0.0102)
5	-0.0109^{***}	$-0.0115*$	-0.0117^{***}	-0.0192^{**}
	(0.0025)	(0.0068)	(0.0022)	(0.0085)
6	-0.0099^{***}	−0.0062	-0.0104^{***}	−0.0119
	(0.0026)	(0.0063)	(0.0024)	(0.0076)
7	-0.0090^{***}	−0.0061	-0.0099^{***}	−0.0106
	(0.0028)	(0.0061)	(0.0026)	(0.0071)
8	-0.0088^{***}	−0.0047	-0.0098^{***}	−0.0090
	(0.0034)	(0.0065)	(0.0032)	(0.0070)
N	1,500,365	1,251,926	1,886,299	1,274,019

Notes: The table reports the effect of the percentage decline in employment at graduation on ln(daily wage). Columns (1) and (2) only include graduates who took the direct track from secondary school from HAVO to University of applied science and from VWO to University. Students who went from VWO to University of applied science or from HAVO to University (usually via VWO or University of applied science) are excluded. Columns (3) and (4) report estimates for workers who are observed as employed for each year since graduation. Coefficients obtained from IV regressions where the percentage decline in employment at graduation is instrumented with the percentage decline in employment in the nominal year of graduation. Regressions include fixed effects for potential experience, calendar year, field of study and graduation year. Demographic controls are age at graduation and gender. They are separately estimated for graduates from universities of applied science and university. Standard errors in parentheses are clustered at the level of graduation cohort and field of study. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
Source: Own calculations based on registration data from Statistics Netherlands.

Table A7: IV estimates using different indicators for match quality.

	University of applied science				University			
	Top 10	Match indicator ≥ 1.5	Top 5 for workers aged 20 - 35	Top 5 with simple shares	Top 10	Match indicator ≥ 1.5	Top 5 for workers aged 20 - 35	Top 5 with simple shares
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Effect at years since graduation								
1	−0.0268***	−0.0083	−0.0170***	−0.0156***	−0.0308***	−0.0322***	−0.0139*	−0.0302***
	(0.0045)	(0.0067)	(0.0039)	(0.0053)	(0.0079)	(0.0105)	(0.0075)	(0.0092)
2	−0.0243***	−0.0139***	−0.0219***	−0.0157***	−0.0322***	−0.0327***	−0.0190**	−0.0301***
	(0.0035)	(0.0040)	(0.0032)	(0.0033)	(0.0078)	(0.0100)	(0.0074)	(0.0092)
3	−0.0214***	−0.0168***	−0.0199***	−0.0154***	−0.0250***	−0.0238***	−0.0139**	−0.0200**
	(0.0032)	(0.0034)	(0.0033)	(0.0027)	(0.0070)	(0.0090)	(0.0070)	(0.0085)
4	−0.0182***	−0.0160***	−0.0177***	−0.0138***	−0.0193***	−0.0160**	−0.0097	−0.0092
	(0.0030)	(0.0032)	(0.0032)	(0.0025)	(0.0063)	(0.0078)	(0.0064)	(0.0076)
5	−0.0166***	−0.0159***	−0.0165***	−0.0139***	−0.0139**	−0.0086	−0.0053	−0.0004
	(0.0029)	(0.0033)	(0.0031)	(0.0026)	(0.0059)	(0.0069)	(0.0061)	(0.0070)
6	−0.0157***	−0.0155***	−0.0156***	−0.0141***	−0.0104*	−0.0030	−0.0025	0.0030
	(0.0030)	(0.0035)	(0.0032)	(0.0027)	(0.0057)	(0.0064)	(0.0060)	(0.0068)
7	−0.0156***	−0.0157***	−0.0153***	−0.0150***	−0.0091	−0.0016	−0.0021	0.0040
	(0.0032)	(0.0038)	(0.0032)	(0.0030)	(0.0057)	(0.0061)	(0.0059)	(0.0068)
8	−0.0166***	−0.0163***	−0.0157***	−0.0149***	−0.0076	0.0001	−0.0006	0.0055
	(0.0035)	(0.0043)	(0.0033)	(0.0033)	(0.0058)	(0.0061)	(0.0059)	(0.0067)
N	2,035,157	2,035,157	2,035,157	2,035,157	1,448,802	1,448,802	1,448,802	1,448,802

Notes: The table reports the effect of the percentage decline in employment at graduation on a dummy for being in a good match. The match measure M_f^s is defined in section 3.2.6. Columns (1) and (5) use the top 10 of sectors instead of the top 5. Columns (2) and (6) use an indicator for whether the match measure is ≥ 1.5. This means that graduates of major m are 1.5 times as likely to work in this sector as the average major. Columns (3) and (7) use the top 5, but only using workers aged 20 - 35 for calculating the match measure. Columns (4) and (8) use the top 5 sectors, but using simple shares of graduates in a sector, rather than the normalized measure. Coefficients obtained from IV regressions where the percentage decline in employment at graduation is instrumented with the percentage decline in employment in the nominal year of graduation. Regressions include fixed effects for potential experience, calendar year, field of study and graduation year. Demographic controls are age at graduation and gender. They are separately estimated for graduates from universities of applied science and university. Standard errors in parentheses are clustered at the level of graduation cohort and field of study. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Source: Own calculations based on registration data from Statistics Netherlands.

A.3 Estimation results for figures in main text

Table A8: Parameter estimates from IV regressions for Figure 3.5.

	University of applied science				University			
	ln(daily wage)	Employment	Self-employment	On benefits	ln(daily wage)	Employment	Self-employment	On benefits
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Effect at years since graduation								
1	-0.0598^{***}	-0.0018^{*}	-0.0004	-0.0013^{**}	-0.0938^{***}	-0.0052^{*}	0.0079^{**}	-0.0040^{***}
	(0.0075)	(0.0010)	(0.0011)	(0.0006)	(0.0177)	(0.0030)	(0.0031)	(0.0012)
2	-0.0277^{***}	-0.0024^{***}	0.0002	-0.0003	-0.0660^{***}	-0.0076^{***}	0.0050^{**}	-0.0014
	(0.0039)	(0.0008)	(0.0008)	(0.0005)	(0.0131)	(0.0028)	(0.0024)	(0.0010)
3	-0.0206^{***}	-0.0018^{**}	0.0006	0.0001	-0.0511^{***}	-0.0044^{*}	0.0018	-0.0002
	(0.0031)	(0.0008)	(0.0008)	(0.0004)	(0.0122)	(0.0023)	(0.0020)	(0.0009)
4	-0.0154^{***}	-0.0017^{**}	0.0007	-0.0000	-0.0351^{***}	-0.0004	-0.0006	0.0000
	(0.0025)	(0.0007)	(0.0007)	(0.0004)	(0.0103)	(0.0021)	(0.0018)	(0.0009)
5	-0.0121^{***}	-0.0012	0.0007	-0.0004	-0.0201^{**}	0.0023	-0.0036^{*}	-0.0008
	(0.0023)	(0.0007)	(0.0007)	(0.0004)	(0.0086)	(0.0022)	(0.0020)	(0.0009)
6	-0.0111^{***}	-0.0008	0.0008	-0.0003	-0.0124	0.0050^{**}	-0.0065^{***}	-0.0015
	(0.0025)	(0.0008)	(0.0007)	(0.0004)	(0.0078)	(0.0023)	(0.0022)	(0.0010)
7	-0.0104^{***}	-0.0005	0.0010	-0.0003	-0.0116	0.0062^{**}	-0.0084^{***}	-0.0010
	(0.0027)	(0.0008)	(0.0008)	(0.0004)	(0.0072)	(0.0025)	(0.0025)	(0.0009)
8	-0.0103^{***}	-0.0007	0.0009	-0.0001	-0.0102	0.0064^{**}	-0.0096^{***}	-0.0016^{*}
	(0.0032)	(0.0010)	(0.0009)	(0.0004)	(0.0073)	(0.0027)	(0.0027)	(0.0010)
N	2,067,535	2,179,304	2,179,304	2,179,304	1,467,864	1,559,875	1,559,875	1,559,875

Notes: The table reports the effect of the percentage decline in employment at graduation on ln(daily wage), and dummies for being employed, being self-employed and receiving benefits. Coefficients obtained from IV regressions where the percentage decline in employment at graduation is instrumented with the percentage decline in employment in the nominal year of graduation. Regressions include fixed effects for potential experience, calendar year, field of study and graduation year. Demographic controls are age at graduation and gender. They are separately estimated for graduates from universities of applied science and university. Standard errors in parentheses are clustered at the level of graduation cohort and field of study. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Source: Own calculations based on registration data from Statistics Netherlands.

Table A9: Parameter estimates from IV regressions for Figure 3.6.

	University of applied science				University			
	Mean firm wage	Firm mobility	1-digit sector mobility	2-digit sector mobility	Mean firm wage	Firm mobility	1-digit sector mobility	2-digit sector mobility
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Effect at years since graduation								
1	-0.0360^{***}				-0.0501^{***}			
	(0.0045)				(0.0112)			
2	-0.0260^{***}	0.0186^{***}	0.0271^{***}	0.0261^{***}	-0.0444^{***}	0.0264^{***}	0.0291^{***}	0.0288^{***}
	(0.0028)	(0.0028)	(0.0027)	(0.0029)	(0.0115)	(0.0053)	(0.0047)	(0.0048)
3	-0.0153^{***}	0.0094^{***}	0.0160^{***}	0.0155^{***}	-0.0304^{***}	0.0133^{***}	0.0204^{***}	0.0198^{***}
	(0.0019)	(0.0020)	(0.0020)	(0.0021)	(0.0102)	(0.0047)	(0.0040)	(0.0041)
4	-0.0075^{***}	0.0058^{***}	0.0107^{***}	0.0110^{***}	-0.0171^{**}	0.0134^{***}	0.0178^{***}	0.0175^{***}
	(0.0016)	(0.0019)	(0.0017)	(0.0018)	(0.0083)	(0.0037)	(0.0035)	(0.0036)
5	-0.0040^{**}	0.0058^{***}	0.0090^{***}	0.0092^{***}	-0.0075	0.0121^{***}	0.0126^{***}	0.0122^{***}
	(0.0016)	(0.0016)	(0.0012)	(0.0013)	(0.0068)	(0.0034)	(0.0030)	(0.0030)
6	-0.0033	0.0033^{**}	0.0057^{***}	0.0063^{***}	-0.0014	0.0047	0.0042	0.0036
	(0.0018)	(0.0016)	(0.0011)	(0.0012)	(0.0061)	(0.0034)	(0.0029)	(0.0030)
7	-0.0035	0.0030^{*}	0.0044^{***}	0.0051^{***}	0.0004	-0.0009	-0.0012	-0.0020
	(0.0019)	(0.0016)	(0.0011)	(0.0012)	(0.0060)	(0.0031)	(0.0028)	(0.0029)
8	-0.0031	0.0010	0.0027^{**}	0.0033^{**}	0.0027	-0.0041	-0.0047	-0.0052^{*}
	(0.0023)	(0.0019)	(0.0013)	(0.0015)	(0.0062)	(0.0034)	(0.0029)	(0.0030)
N	2,035,157	1,924,495	1,924,495	1,924,495	1,448,802	1,368,067	1,368,067	1,368,067

Notes: The table reports the effect of the percentage decline in employment at graduation on ln(mean firm wage) and the probability to switch firm, 1-digit or 2-digit sector. Coefficients obtained from IV regressions where the percentage decline in employment at graduation is instrumented with the percentage decline in employment in the nominal year of graduation. Regressions include fixed effects for potential experience, calendar year, field of study and graduation year. Demographic controls are age at graduation and gender. They are separately estimated for graduates from universities of applied science and university. Standard errors in parentheses are clustered at the level of graduation cohort and field of study. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
Source: Own calculations based on registration data from Statistics Netherlands.

Table A10: Parameter estimates from IV regressions for Figure 3.7.

	University of applied science				University			
	Good match	Low-ranked firm	Medium-ranked firm	High-ranked firm	Good match	Low-ranked firm	Medium-ranked firm	High-ranked firm
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Effect at years since graduation								
1	-0.0187^{***}	0.0097^{***}	0.0117^{***}	-0.0150^{***}	-0.0168^{**}	0.0156^{***}	0.0095	-0.0232^{***}
	(0.0034)	(0.0016)	(0.0026)	(0.0028)	(0.0074)	(0.0036)	(0.0071)	(0.0083)
2	-0.0212^{***}	0.0086^{***}	0.0125^{***}	-0.0138^{***}	-0.0206^{***}	0.0136^{***}	0.0131^{**}	-0.0253^{***}
	(0.0031)	(0.0015)	(0.0022)	(0.0025)	(0.0073)	(0.0034)	(0.0065)	(0.0085)
3	-0.0188^{***}	0.0081^{***}	0.0094^{***}	-0.0090^{***}	-0.0162^{**}	0.0096^{***}	0.0182^{***}	-0.0227^{***}
	(0.0032)	(0.0014)	(0.0021)	(0.0023)	(0.0069)	(0.0031)	(0.0058)	(0.0076)
4	-0.0158^{***}	0.0072^{***}	0.0062^{***}	-0.0055^{**}	-0.0116^{*}	0.0072^{***}	0.0164^{***}	-0.0160^{**}
	(0.0031)	(0.0015)	(0.0020)	(0.0024)	(0.0063)	(0.0027)	(0.0050)	(0.0065)
5	-0.0145^{***}	0.0060^{***}	0.0057^{***}	-0.0055^{**}	-0.0071	0.0061^{**}	0.0147^{***}	-0.0115^{**}
	(0.0031)	(0.0015)	(0.0018)	(0.0024)	(0.0059)	(0.0024)	(0.0046)	(0.0056)
6	-0.0135^{***}	0.0052^{***}	0.0048^{**}	-0.0053^{*}	-0.0048	0.0041^{*}	0.0134^{***}	-0.0072
	(0.0031)	(0.0015)	(0.0019)	(0.0027)	(0.0058)	(0.0022)	(0.0044)	(0.0051)
7	-0.0135^{***}	0.0048^{***}	0.0044^{**}	-0.0057^{**}	-0.0042	0.0040^{*}	0.0106^{**}	-0.0048
	(0.0031)	(0.0014)	(0.0017)	(0.0027)	(0.0057)	(0.0022)	(0.0044)	(0.0050)
8	-0.0137^{***}	0.0043^{***}	0.0029	-0.0049	-0.0033	0.0034	0.0095^{**}	-0.0036
	(0.0032)	(0.0015)	(0.0019)	(0.0031)	(0.0057)	(0.0023)	(0.0046)	(0.0053)
N	2,035,157	2,179,304	2,179,304	2,179,304	1,448,802	1,559,875	1,559,875	1,559,875

Notes: The table reports the effect of the percentage decline in employment at graduation on a dummy for match quality and dummies for being in a low-, medium- and high-ranked firm. Coefficients obtained from IV regressions where the percentage decline in employment at graduation is instrumented with the percentage decline in employment in the nominal year of graduation. Regressions include fixed effects for potential experience, calendar year, field of study and graduation year. Demographic controls are age at graduation and gender. They are separately estimated for graduates from universities of applied science and university. Standard errors in parentheses are clustered at the level of graduation cohort and field of study. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
Source: Own calculations based on registration data from Statistics Netherlands.

A.4 Descriptives on each field of study

Table A11: Descriptive statistics on the change in employment at graduation for each field of study.

		Change in employment at graduation				
ISCED 2011	Label	Mean	Std. Dev.	Min	Max	Observations
University of applied science						
14	Teacher training and education science	1.30	1.23	−0.93	3.19	59,060
21	Arts	1.40	1.52	−0.98	4.20	12,426
22	Humanities	1.33	1.44	−1.07	3.77	1,086
31	Social and behavioural science	1.09	1.79	−1.62	3.73	11,873
32	Journalism and information	0.86	1.56	−1.52	4.32	6,188
34	Business and administration	1.10	2.05	−1.89	4.37	64,798
38	Law	1.11	1.49	−0.82	3.46	1,058
42 - 46	Science and Mathematics	1.55	1.84	−0.91	4.79	697
48	Computing	3.22	4.47	−1.95	13.94	13,531
52 - 54	Engineering and manufacturing	1.15	2.31	−2.53	4.98	18,361
58	Architecture and building	0.79	1.87	−2.69	3.22	11,903
62 - 64	Agriculture and veterinary	0.86	1.49	−1.32	3.09	4,585
72	Health	2.07	1.48	−0.52	3.99	37,916
76	Social services	1.72	1.72	−1.81	4.29	26,455
81	Personal services	1.35	1.91	−1.28	4.46	22,910
84	Transport services	1.09	1.93	−2.27	4.06	4,049
85 - 86	Environmental protection and security services	0.98	1.32	−0.82	3.15	2,050
University						
14	Teacher training and education science	1.37	1.29	−0.91	3.58	14,464
21	Arts	1.44	1.64	−1.09	4.17	5,209
22	Humanities	1.26	1.38	−0.98	3.76	13,276
31 - 32	Social and behavioural science	1.35	1.72	−0.95	4.01	52,677
34	Business and administration	1.27	2.01	−1.84	4.42	37,115
38	Law	1.07	1.64	−1.23	3.51	24,329
42	Life science	1.60	1.48	−0.68	4.12	3,754
44	Physical science	1.52	1.97	−1.45	5.17	5,183
46	Mathematics and statistics	1.62	2.36	−1.12	5.97	1,248
48	Computing	2.99	4.13	−1.93	12.80	5,499
52 - 54	Engineering and manufacturing	1.34	2.37	−2.59	5.55	8,605
58	Architecture and building	1.28	1.91	−2.55	3.77	7,776
62 - 64	Agriculture and veterinary	1.09	1.38	−1.03	3.46	2,824
72	Health and social services	2.05	1.33	−0.15	4.16	31,301
81, 84 - 86	Services	1.33	1.48	−0.71	3.34	2,794

Notes: The number of observations in the final column refers to the total number of graduates in that field in our sample.
Source: Own calculations based on registration data from Statistics Netherlands.

A.5 Descriptives on firm rank, job mobility and match quality

Table A12: Descriptive statistics on job mobility.

	University of applied science			University		
	Firm mobility	1-digit sector mobility	2-digit sector mobility	Firm mobility	1-digit sector mobility	2-digit sector mobility
	(1)	(2)	(3)	(4)	(5)	(6)
Probability at years since graduation						
2	40.75	25.98	28.32	41.95	28.47	30.37
3	31.88	17.98	19.77	32.80	18.56	20.06
4	26.84	14.14	15.76	28.68	15.52	16.89
5	24.72	12.59	14.10	26.44	13.71	15.09
6	22.32	10.86	12.27	25.35	12.82	14.12
7	20.53	9.59	10.89	23.34	11.27	12.51
8	19.27	8.66	9.92	21.61	9.94	11.12

Notes: The table reports the probability of switching firm, 1-digit and 2-digit sector at each year since graduation.
Source: Own calculations based on registration data from Statistics Netherlands.

Table A13: Descriptive statistics on firm rank.

	University of applied science				University	
	Low-ranked firm	Medium-ranked firm	High-ranked firm	Firm mobility	1-digit sector mobility	2-digit sector mobility
	(1)	(2)	(3)	(4)	(5)	(6)
Probability at years since graduation						
1	20.08	48.49	31.43	15.50	40.44	44.06
2	16.78	46.66	36.56	11.85	37.03	51.12
3	15.19	44.75	40.06	10.26	33.08	56.66
4	14.24	43.54	42.22	9.12	30.95	59.93
5	13.60	42.79	43.61	8.40	29.77	61.83
6	12.99	42.18	44.83	7.69	29.60	62.71
7	12.46	41.88	45.67	7.40	29.17	63.43
8	12.13	41.47	46.40	7.05	28.96	63.99

Notes: The table reports the probability of working in a low-, medium- or high-ranked firm at each year since graduation.
Source: Own calculations based on registration data from Statistics Netherlands.

CHAPTER 4

Do Parents Work More When Children Start School? Evidence from the Netherlands*

4.1 Introduction

The first day a child enters school is a very exciting day in his or her life: it marks the start of many new experiences and the acquisition of many new skills. As economists put it: children formally start investing in their human capital, and this process will continue to influence their lives for many years to come. The benefits of going to school on child development have been studied extensively: compulsory schooling does not only improve test scores, but also causally enhances later life outcomes, such as wages and health (Devereux and Hart, 2010; Machin et al., 2011; Grenet, 2013).

At the same time, compulsory schooling also provides benefits for parents, as they start saving time and/or money – at least in countries where the government subsidizes primary schools more than child care – which may affect their labor supply. On the one hand, parents who used to take care of their children during school hours are expected to increase their labor supply when their youngest child starts school. These parents do not need to take care of their children anymore and hence have free time on their hands. On the other hand, parents whose children

*This chapter is joint work with Lisette Swart and Karen van der Wiel. It is based on Swart et al. (2019).

attended (paid) childcare before going to school might decrease their labor supply, when their youngest child starts school. These households save on childcare expenses and therefore experience an income effect.

In this paper, we empirically estimate the effect of children starting compulsory schooling in the Netherlands on their parents' labor market position. Using a balanced panel of administrative data on all Dutch parents between 2006 and 2016, we analyze the labor market position of parents when their youngest child is between three and six years old, and we do this for mothers and fathers separately. We apply a difference-in-differences approach to tease out any changes observed due to macro-level shocks, increased working experience over these three years or other changes unrelated to school. The control group consists of parents with a youngest child between one and four years old. They therefore do not experience the 'treatment' because their youngest child does not start school during the observation period. To ensure that our control and treatment group are as similar as possible and to take cohort effects into account, we apply coarsened exact matching.

How compulsory schooling affects the parents' labor market position is of international interest as employment of mothers of young children is relatively low in most countries, and policy makers are hence interested in the drivers of their labor supply. Interestingly, Figure 4.1 shows that in most European countries the differences in the average number of hours worked by mothers with or without all their children in school are small. It is unclear however to what extent this is driven by cohort effects: younger cohorts with smaller children tend to work more in general.

The Netherlands is a particularly interesting country to study labor market effects for parents of compulsory schooling. Primary education is paid for by the government, while child care is not, and child care subsidies are dependent on household income. Moreover, children start school (kindergarten) for approximately 20 hours a week in the month that they turn four, which makes it possible to rule out seasonal employment effects. Furthermore, the participation of Dutch women is relatively high, while their working hours are relatively low compared to that of women in other European countries as the majority of Dutch women work part-time (see Figure 4.1). This means that there is substantial room for increases in hours worked. A large majority of the men however, works fulltime: in 2016, the employment rate among men was 82.6 percent and more than three quarters works fulltime.

Our findings show economically small labor supply reactions and that mothers adapt their working hours significantly more than fathers when their youngest child starts school. On average, households save approximately 60 to 70 euros per month

Figure 4.1: Hours worked by mothers before and after the youngest child in a household starts going to school

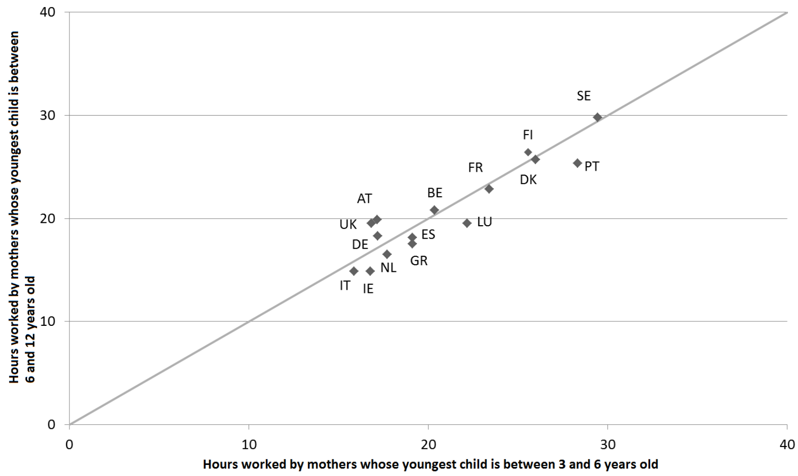

Source: Own calculations based on Eurostat (2014).
Notes: The x-axis indicates how many hours women work when the youngest child in their household is older than three years old, but does not yet go to school and the y-axis denotes the number of hours women work when the youngest child in their household goes to school and is younger than twelve years old. These averages take all mothers in a country into account: mothers who do not work, work zero hours. In addition, this figure is based on cross-sectional data, so that differences may also be affected by cohort effects: overall labor participation of younger generations of women exceeds that of older generations. Finally, the solid line denotes the diagonal: countries in which mothers of children who do not go to school work as much as mothers whose children do go to school, are on the diagonal. In countries above (below) the diagonal, women whose youngest child starts going to school work longer (shorter) hours.

on child care. Dutch mothers experience an increase in their free time of 13 hours a week when their youngest child goes to school. After two years, we find that the average number of hours Dutch mothers work increases by around 0.5 hours. These changes are driven for about two-thirds by responses at the intensive and one-thirds at the extensive margin. Given that the average mother in our sample works around 15 hours a week, mothers increase their work hours by 3% after their youngest child goes to school. Empirically, the income effect is thus dominated by the effect of additional time on labor supply. Dutch fathers experience an increase in their free time of almost 4 hours a week when their youngest child goes to school. We find that they increase working hours on average by around 0.15 hours, or 0.4% relative to the mean. This is for about two-thirds driven by a response along the extensive margin and one-thirds at the intensive margin.

Heterogeneity analyses suggest that the response is larger among mothers and fathers who are already working longer hours and who already earn higher wages. This is perhaps surprising because they experience smaller time gains and they typically save more money on child care expenses when their youngest child starts school. Theoretically, one would thus expect smaller increases in working hours in these groups. However, differences in initial preferences for labor supply over leisure are probably at the heart of these results: these parents are more likely to have decreased working hours solely to be with their children and hence return to the office once child care activities are less needed. Single mothers show an overall negative employment response to their youngest child attending school. The income effect apparently dominates for this subgroup. One explanation for the relatively large income effect is mental accounting (Thaler, 1985; Abeler and Marklein, 2016). In the Netherlands, child care subsidies depend on household income, are paid by the tax authorities, and parents generally receive them on a different date than when they pay the child care institution. It is therefore possible that parents do not take the net monetary gain they obtain when their youngest child starts going to school into account, but rather consider the gross monetary gain of reduced expenses on child care without taking the corresponding decline in child care subsidies into account. This could lead to much larger perceived income gains. The difference between the gross and the net monetary gain is generally substantial; the gross monetary gain can be up ten times as large as the net monetary gain for low-income households.

Our analysis makes several contributions to the literature on compulsory schooling and the labor supply of parents, which generally uses data on earlier cohorts of parents. The first contribution of this paper is that findings are not confounded by seasonal effects in labor supply. In the Dutch institutional setting, children go to school on the day they turn four years old. This provides us with variation in school entry over the calendar year, unlike other studies from countries where children always start school in September (Gelbach, 2002; Goux and Maurin, 2010; Finseraas et al., 2016). The second contribution is that we can precisely estimate the effects using very similar treatment and control groups. This is because we use a recent administrative dataset, which includes the hours a recent cohort of parents work in each month as well as information on child care subsidies for the universe of Dutch parents. In addition, the data also enable us to determine the heterogeneity in the magnitude of the shock parents experience in terms of money and time depending on their working hours and the number of hours their youngest child attends formal child care. The third contribution is that we symmetrically estimate effects for both mothers and fathers, while most other papers in the literature focus on the employment effects of children going to school for mothers only.

Interestingly, our heterogeneity analyses provide different results than most of the literature; often older papers find stronger effects for single women and low-income groups, but we draw opposite conclusions. A paper by Gelbach (2002) for example examines the effect of public schooling for five year olds in the US on their mothers' labor supply and finds that single mothers increase the number of hours they work by 6-24 percent, while for married mothers he finds an increase of 6-15 percent. Gelbach uses cross-sectional data from the 1980 US Census, while we use panel data from 2006 to 2016. Goux and Maurin (2010) analyze census data from 1999 in France where all children start school at age 3 or earlier using a cut-off within the year for eligibility and find that pre-elementary school only has a significant effect on the labor supply of single mothers. Similarly, Finseraas et al. (2016) examine a reform in Norway in 1997 where the compulsory school starting age was lowered from six to five. They find a short-term increase of labor supply, with much stronger effects for mothers with low wage potential, who probably did not use childcare. Recent cohorts differ in important dimensions from older cohorts: they are higher educated, employment rates have increased steadily (Blundell et al., 2013) and labor supply elasticities are currently smaller (Blau and Kahn, 2007).

Another related branch of the literature looks at the effect of the availability of – non-compulsory – kindergarten on the labor market position of parents. Cascio (2009) exploits the introduction of subsidies for kindergarten in school districts in the US from the 1960s to the 1980s to examine effects of public schooling on labor supply of mothers. Using cross-section data from the Census for each decade from 1950 to 1990 she only finds significant effects for single mothers. Fitzpatrick (2010) considers pre-kindergarten availability in three states in the US. Using Census data from 2000 in an RD design, she finds no robust impact of pre-kindergarten availability on maternal labor supply. Fitzpatrick (2012) uses the same data, but exploits cutoffs for eligibility for public kindergarten. She finds that eligibility only increases the employment of single mothers without additional children. Related, Shure (2019) considers the effect of the extension of the schoolday in Germany in the 2000s on maternal labor supply. She finds no effect along the intensive margin, but an increase in employment probability.

4.2 Theoretical Framework

This section formalizes the intuition of how parental labor supply is affected when children start going to school into testable hypotheses. To do this, we analyze the changes in the budget constraint parents face that occur when their youngest child starts going to school. For simplification, we restrict our focus to three activities parents can spend their working hours on[1]: (1) working on the labor market l_i, (2) taking care of children at home t_i or (3) leisure z_i. The time constraint for each individual i is therefore:

$$l_i + t_i + z_i = 1 \quad \text{where} \quad 0 \geq l_i, t_i, z_i \geq 1 \tag{4.1}$$

Children always have to be taken care of. This can either be done by the father t_f or the mother t_m, or by child care providers t_y. Older children also spend a number of hours at school, t_s. As long as at least one child in a household does not go to school yet, parents have to arrange child care during school hours. The time constraint of taking care of the children can therefore be characterized by:

$$t_f + t_m + t_y + t_s = 1 \quad \text{where} \quad t_s = 0 \quad \text{if at least one child does not go to school} \tag{4.2}$$

$$0 < t_s < 1 \quad \text{if all children go to school.}$$

Households earn income from labor l_i and spend money on child care t_y which costs q and on consumption X which costs p. Households therefore face the following budget constraint:

$$qt_y + pX \leq w_f l_f + w_m l_m \tag{4.3}$$

In this budget constraint people cannot save any money for future consumption. Substituting the time constraints for fathers and mothers (Equation 4.1) as well as the time constraint for child care (Equation 4.2) into this equation yields the following budget constraint:

$$q(1 - t_m - t_f - t_s) + pX \leq w_f(1 - t_f - z_f) + w_m(1 - t_m - z_m) \tag{4.4}$$

When the youngest child in a household starts going to school, the budget constraint parents face changes. Figure 4.2 graphically illustrates the changes one parent faces, keeping the time the other parent spends on labor l_f and child care t_f constant. In the Figure, we show the mother's budget constraint, but the budget constraints fathers face keeping the time mothers spend on labor and child care constant is identical. Mothers decide how to allocate their time between working which enables them to consume X, leisure z_m and taking care of their children

[1] The time available during the (work) week is normalized to 1 without loss of generality. Outside of working hours parents can still spend time with their children and enjoy leisure.

t_m. This three-dimensional budget constraint mothers face is shown in three two-dimensional views in Figure 4.2. The blue solid planes represent the situation when at least one child does not go to school yet and the green dashed planes represent the situation after the youngest child has started going to school and parents no longer have to arrange child care during school hours.

Figure 4.2: The mother's budget constraint.

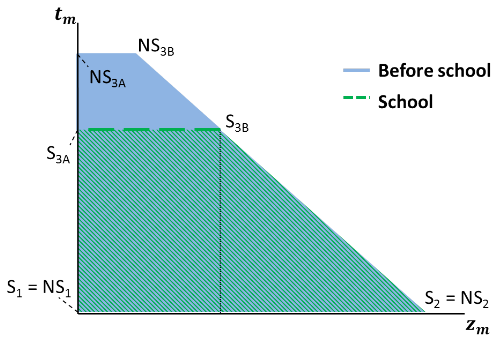

(a) Domestic child care t_m vs leisure z_m

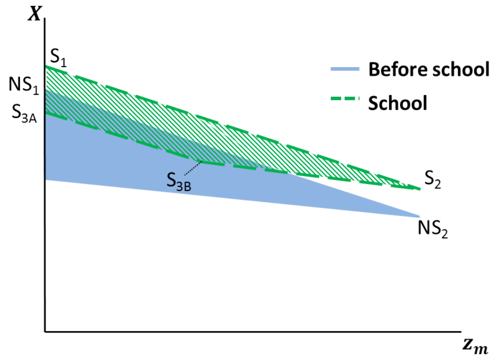

(b) Consumption X vs leisure z_m

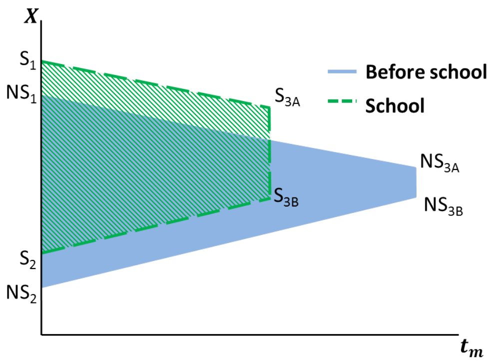

(c) Consumption X vs child care t_m

Legend

Situation	l_m	z_m	t_m	X
1	1	0	0	$w_f^* l_f + w_m^* - q^*(1 - t_s - t_f)$
2	0	1	0	$w_f^* l_f - q^*(1 - t_s - t_f)$
3A	$t_s + t_f$	0	$1 - t_s - t_f$	$w_f^* l_f + w_m^*(t_s + t_f)$
3B	0	$t_s + t_f$	$1 - t_s - t_f$	$w_f^* l_f$

Notes: This Figure shows three two-dimensional views of the household budget constraint that parents face when deciding on how to allocate their time. The blue solid planes represent the situation before the youngest child goes to school, where $t_s = 0$, while the green dashed planes denote the situation after the youngest child started school, $t_s > 0$. In the figures, situations before the child starts going to school are denoted by NS_i and situations once all children in a household go to school are referred to as S_i. The numbers in the subscript refer to situations shown in the legend. Appendix Section A.1 explains in more detail how this budget constraint changes when the youngest child in a household starts going to school. Although the highest point of the budget constraint before children go to school (the blue solid plane) in Figure 4.2c is higher than the kink of the budget constraint after children go to school (the green dashed plane), these relative positions could be reversed without affecting the hypotheses.

When the youngest child goes to school, parents do not need to pay for child care for the hours during which the children are at school. Parents who pay for child care therefore experience an income effect. This income effect shifts the budget constraint upward and causes mothers to decrease their labor supply. At the same time, the maximum amount of time parents can spend on domestic child care reduces, because children do not need to be taken care of by their parents when they are at school. This 'time effect' causes parents who used to take care of their children during school hours prior to the youngest child going to school, to increase their labor supply. This therefore yields the following hypotheses, which are not mutually exclusive:

Hypothesis 1 *When the youngest child in the household starts going to school, working mothers experience an income effect, which decreases their labor supply.*

Hypothesis 2 *When the youngest child in the household starts going to school, mothers who do not work fulltime obtain additional time, which increases their labor supply.*

This theoretical framework also allows for the prediction of some heterogeneous effects. If the real wage of the mother increases ceteris paribus, she experiences both a substitution and an income effect. The income effect decreases labor supply, while the substitution effect has an upward effect. Generally, the substitution effect dominates the income effect. For higher real wages of the mother, the negative effect on maternal labor supply is therefore smaller. Additionally, if the price of child care decreases ceteris paribus, the marginal benefit of working increases, making it less attractive to decrease labor supply.

In sum, when the youngest child starts going to school, this may affect maternal labor supply in two ways. On the one hand, parents whose children attend external child care experience an income effect that may reduce labor supply. On the other hand, parents who provide domestic child care spend less time doing so which may increase their labor supply.

4.3 Magnitude of the treatment

This section discusses the changes parents experience when the youngest child in a household starts going to school. The magnitude of these changes determines the treatment parents experience. This magnitude depends on the number of hours parents work, the number of hours children attend formal child care and on the price of child care. We start this section by explaining the institutional setting in the Netherlands concerning school and child care. Subsequently, we calculate the (counterfactual) magnitude of the change mothers and fathers experience in terms of time when their children start going to school, followed by a calculation of the magnitude of the change households experience in terms of money. To calculate these shocks, we determine how large the change would have been if the parental labor supply and the number of hours children attend external child care had remained the same.

4.3.1 Institutional setting

In the Netherlands, children generally start school the day they turn four years old. Although schooling is only compulsory from the age of five, more than 99 percent of children starts school when they are four years old (Eurostat, 2012). In the first years of primary school, children attend school for on average 22 hours a week, or on average for 4.4 hours a day.[2]

Formal child care in the Netherlands is privatized and institutions can set their own price.[3] Parents can obtain a child care subsidy from the government for the hours children between the age of 0 and 12 attend formal child care.[4] This subsidy is a percentage of the hourly price child care institutions charge. The subsidy varies with household income: the government pays a lower share of the costs of child care for households with a higher income than for households with a lower income. Primary school on the other hand is paid for by the government, although schools may ask for a 'voluntary parental contribution' which may range from 10 euros a year to 120 euros a year.

4.3.2 Magnitude of the change in terms of time

This subsection calculates the number of hours parents save when their youngest child starts going to school. Parents only save time when they do not work fulltime

[2]The number of hours children go to school varies between schools, because the only requirement set by the government is that in the first four years of primary school, children need to receive education for at least 3,520 hours (Rijksoverheid, 2017). This is 22 hours a week on average.

[3]The government annually sets a threshold price per hour and only for the amount parents pay below this threshold they receive government subsidies.

[4]Formal child care includes day care for children younger than four years old, out-of-school care for children in primary school and childminders for children between 0 and 12 years old.

and provide domestic child care during working hours. When the youngest child in a family starts going to school, parents no longer need to take care of the children during school hours.

Fulltime employment in the Netherlands takes up forty hours a week. A regular working day therefore consists of eight hours and children go to school for on average 4.4 hours a day in the first years of primary school. The shock parents experience in terms of time therefore consists of 4.4 hours a day for each day that a parent does not work, in other words:

$$\text{Time shock } = (40 - \text{hours worked}) \times \frac{4.4}{8} \tag{4.5}$$

In these calculations, we assume that working hours are fixed to eight hours a day. This implies that our calculation of the time shock is in fact an upper bound. Parents who have flexible working hours may experience a shock smaller than the one calculated here.

On average mothers save a substantial 13.3 hours per week when their youngest child starts going to school, as illustrated in Column 2 of Panel A in Table 4.1. Column 2 of Panel B illustrates that the shock fathers experience in terms of time is substantially smaller, only 3.7 hours on average. Fathers generally save less time because a much larger share of them works fulltime, as is illustrated in Column 1 of Table 4.1. Parents who work fulltime hardly save any time when their youngest child goes to school, as they generally do not take care of the children during working hours before the children start going to school. The time shock is largest for parents who work fewer hours.

4.3.3 Magnitude of the change in terms of money

Parents may also save money when their children start going to school. This only holds true for parents whose children attend child care, as they do not need to pay for child care during the hours the child is at school. When children start school, the number of hours they attend formal child care therefore reduces by approximately fifty percent.[5] The net income-shock parents experience due to reduced expenditures on formal child care can therefore be approximated using the following formula:

$$\text{Monetary shock } = p_N q_N - p_O q_O = (sP_N)(50\% q_O) - (sP_O)q_O = sq_O\,(0.5P_N - P_O) \tag{4.6}$$

[5] Child care providers generally charge parents for ten hours per day for children who have not yet started going to school yet and for five hours per day for out-of-school care for children who attend primary school.

Table 4.1: Magnitude of the shock in terms of money and in terms of time

Panel A: Shock in terms of time and money that mothers experience when their youngest child goes to school	(1) % of mothers	(2) Hours saved per week	(3) Money saved per month (per household)
Does not work	25.6	22.0	8
Works less than 20 hours a week	30.0	14.9	39
Works 20-34 hours a week	39.6	7.9	103
Works more than 34 hours a week	4.7	1.5	135
Total	100.0	13.3	61

Panel B: Shock in terms of time and money that fathers experience when their youngest child goes to school	% of fathers	Hours saved per week	Money saved per month (per household)
Does not work	6.4	22.0	22
Works less than 20 hours a week	4.3	16.1	43
Works 20-34 hours a week	16.6	5.2	83
Works more than 34 hours a week	72.7	1.0	72
Total	100.0	3.7	70

Notes: See main text for calculation of the monetary and time shocks. The percentages listed in Column (1) refer to the share of mothers (fathers) in our sample that works a specific number of hours. The amounts of money households save when their youngest child starts going to school (listed in Column (3)) differs for the mothers (Panel A) and fathers (Panel B) in our sample, because households with single mothers *are* included in the sample for mothers, but not in the sample of fathers.

where subscript O denotes the old situation (day care) and N denotes the new situation (out-of-school care). Additionally, p denotes the price parents pay for an hour of formal child care, which is a share s of the actual hourly price of child care P: $p = sP$.[6] The price of out-of-school care p_N may differ from the price of day care p_O.[7] Finally, q refers to the number of hours a child attends formal child care, where

[6]For households where both parents work or are enrolled in education, the government subsidizes formal child care, which includes day care, out-of-school care and childminders for children between 0 and 12 years old. The subsidy is a percentage of the price that depends on household income. The share paid by the government varied substantially over the past decade. Formal child care in the Netherlands is privatized and institutions can set their own price. The government annually sets a threshold price per hour set and only for the amount parents pay below this threshold they receive government subsidies.

[7]For these hourly prices of formal child care, we use the average price reported by the sector association of child care providers in the Netherlands in a specific year (Kinderopvang, 2016). The share paid by parents is calculated using a conversion table from the Dutch government that specifies this share for each calendar year for each level of household income.

the number of hours in out-of school care q_N is half the number of hours in the day care, $q_N = 0.5q_O$.[8]

On average, households save approximately 60 to 70 euros a month on the costs of formal child care when their youngest child starts going to school, as shown in Column 3 of Table 4.1. Panel A shows that households where mothers work more hours save more money in absolute terms than households where mothers work fewer hours per week, because household income is generally higher - and childcare subsidies are therefore lower - in households where mothers work more hours.[9] Panel B shows that for fathers this relationship is not linear: households where the fathers work 20-34 hours save more money than households where fathers work more, because on average women work more hours in households where fathers work 20-34 hours.

It is important to note that these average changes in the money parents spend on child care when their children go to school are a lower bound the true magnitude of the monetary shock. Our calculation only takes the costs of *formal* child care into account, as there are no data for all parents on the use of informal child care. If parents use paid informal child care, the true magnitude of the monetary shock exceeds our calculations.

4.4 Empirical Approach

To determine the effect of children going to school on parental labor supply, we employ a difference-in-differences design. We construct a control group of parents with children aged one to four to control for general business cycle effects, policy changes and other time-changing variables. To ensure comparability of the treatment and control group, we employ coarsened exact matching and we do the analyses separately for mothers and for fathers.

[8] The hours a child attends formal child care are available per calendar year. We therefore use the number of hours in the calendar year in which the child turns three years old, i.e. the calendar year before the child starts going to school.

[9] To be eligible for child care subsidies in the Netherlands, both parents need to work or be enrolled in education. Nonetheless, Table 4.1 shows that women who do not work (those working zero hours a week) on average still save some money when their children start going to school. This is likely due to mothers who are enrolled in education or mothers who recently became unemployed and who are therefore still eligible for child care subsidies. Furthermore, before the reform of 2011, the eligibility for subsidies for formal child care did not depend on the number of hours parents worked. Therefore, before 2011 children of women who did not work could also attend formal child care so that these women also experienced a small monetary shock the moment their youngest child started primary school.

4.4.1 Defining treatment and control group

Our treatment group consists of parents whose youngest child is between three and six years old. Since children go to school when they are four years old, this means that a parent is in our sample one year before the youngest child goes to school until two years after the child goes to school. To ensure that parents do in fact experience a time shock when their youngest child goes to school, we impose that the child remains the youngest child in the family until they are six years old. We end up with a balanced sample that consists only of parents who took care of this particular child for each of the three years in the sample.

To estimate the effect of the youngest child going to school on parental labor supply, we compare this treatment group to a control group consisting of parents whose youngest child is between one and four years old. In this group the youngest child does not start going to school yet during the observation period, so that parents do not experience the treatment.[10] With this design, we compare parents whose youngest child turns four years old in a given year (the treatment group) with parents whose youngest child turns two years old in that same year (the control group).

In the robustness checks we also present results where the control group consists of parents whose second-youngest child is between three and six years old. This group does not experience a time shock, since they still have a younger child at home, but they do experience a small monetary shock, since they pay for fewer hours of child care for their second-youngest child when this child starts going to school.[11]

4.4.2 A difference-in-differences design

We estimate the effect of children going to school on the number of hours worked by the parent in a certain month t (those who do not work, work zero hours). We then check whether this effect is driven by changes along the intensive or the extensive

[10] Older children of parents in the control group may still start going to school in the observation period. This does not cause a time shock for the parents, since they still have a younger child at home. Parents do experience a small monetary shock when older children start going to school, since they pay for fewer hours of child care for these children during school hours. This monetary shock is however substantially smaller than the shock parents experience when the youngest child starts going to school (see footnote 9). As a robustness check, we also restricted our control group to parents whose second-youngest child is at least three years older to make sure no child in the household starts going to school within the observation period. This restriction severely limits the sample, but does not affect our results.

[11] This monetary shock is however substantially smaller than the shock parents experience when the youngest child starts going to school, since own contributions for child care are generally highest for the youngest child. The magnitude of the shock parents experience when their second-youngest child starts going to school is shown in Appendix Table A1.

margin by estimating the effect of children going to school on i) a dummy variable that indicates whether the parent works in month t, and ii) the number of hours worked by the parent in month t, excluding those who do not work over the entire observation period. We estimate the following linear model

$$y_{it} = \alpha + \beta \times T_i + \sum_{t=-10}^{24} \gamma_t \times S_t + \sum_{t=-10}^{24} \delta_t \times S_t \times T_i + \eta_i + \theta_t + \varepsilon_{it}, \tag{4.7}$$

where y_{it} is some measure of labor market participation (hours worked or participation). T_i is a dummy variable that takes value 1 if parent i is in the treatment group and 0 if the parent is in the control group. S_t is a set of dummy variables for each month relative to the treatment period, where $S_t = 0$ when the youngest child is four years old (or two years old in the control group). Hence, γ_t captures the general time trend. The coefficients of interest are δ_t, which capture the average effect of the youngest child going to school for the treatment group (ATT) for each month relative to treatment. Our reference category is 11 months before treatment. We include individual fixed effects η_i to control for any individual-specific effects (e.g. the number of children in the household, education level of the parents or ability) that might drive changes in our outcome variable. We also include a full set of calendar year by month fixed effects θ_t to control for any common time shocks (e.g. a recession). Finally, ε_{it} is the error term and standard errors are clustered at the level of the parent to take into account within-parent correlation in labor market behavior.

We also estimate a simpler version of equation 5.7 to get an average treatment effect over the entire two-year post treatment period

$$y_{it} = \alpha + \beta \times T_i + \gamma_t \times S_t + \delta \times S_t \times T_i + \eta_i + \theta_t + \varepsilon_{it}, \tag{4.8}$$

where S_t is a dummy equal to one if the youngest child is in school and and zero otherwise and T_i is a dummy equal to one if the parent is in the treatment group and zero otherwise. The average treatment effect on the treated is given by δ. All other terms are defined as before. In the heterogeneity analyses we interact the treatment dummy with indicators for different groups (e.g. marital status or number of children).

To ensure that we compare parents in the treatment and the control groups who are as similar as possible, we apply coarsened exact matching (CEM) (Azoulay et al., 2010; Iacus et al., 2012b).[12] Following Azoulay et al. (2010), we match on the pre-treatment outcomes. In particular, we match on hours worked by a parent 11 and 6 months before the treatment and on the calendar year of treatment. The hours a parent worked are divided in bins.[13] In this way, we obtain weights for the control groups that are proportional to the number of treated and controls in each combination of (bins of) hours worked 11 and 6 months before treatment. All treated parents can be matched to at least one parent in the control group. Only 0.0001% of the potential control group is not matched.[14]

This matching procedure alleviates concerns regarding the common trend assumption, which is required in our difference-in-differences design. That is, conditional on fixed effects and observables, trends in the treatment and control group should follow a similar trend in the absence of treatment. The matching procedure we apply should reduce possible differences in the pre-trend period. We also explicitly test for common trends in outcomes during the pre-treatment period by estimating the 11 months before treatment.

4.4.3 Data and descriptive statistics

We use administrative data on the universe of employees in the Netherlands.[15] The data consist of monthly employment records collected for the purpose of social security administration and income taxes between 2006 and 2016. The data contain information on actual hours worked and the wage for all jobs of a worker in that month.[16] From these administrative data, we construct a balanced panel of monthly data on parents' average weekly working hours.[17] We add data from the municipal registries (*Gemeentelijke Basisadministratie*) which contains demographic information

[12]In contrast to propensity score matching, CEM is nonparametric. CEM involves selecting covariates on which we want to achieve balance, and then match each treated observation to control observations based on the values of the covariates. The approach is coarse in the sense that we do not match on each value for each covariate, but rather coarsen the distribution for some covariates and achieve balance on the coarsened distribution.

[13]The bins are defined in the following way: 0, 1-4, 5-9, 10-14, 15-19, 20-24, 25-29, 30-34, 35-39 and 40 hours and above.

[14]Our results are comparable if we do not apply CEM.

[15]The data are available to researchers who sign a confidentiality agreement with Statistics Netherlands. While we cannot share the data, all programs to replicate our results are available on request.

[16]We use actual hours worked rather than contractual hours, since actual hours worked also take parental leave into account. Parental leave policies differ substantially between sectors in the Netherlands.

[17]Using the monthly data, we construct average hours worked per week by dividing the number of hours worked in a month by 4.35 (the average number of weeks per month).

on all households, including for each household member the month and year of birth and gender.

We construct separate samples for mothers and their partners, who are the biological fathers or the stepfathers to the children, independent of whether they are married or not.[18] We therefore do include single mothers in our sample, but not single fathers.

Subsequently we apply the following sample restrictions. First, we drop mothers who have their first child before the age of fifteen or after the age of 40 (1%) as well as fathers who have their first child before 15 or after 50 (1%), parents who work more than 100 hours per week during at least one month (0.02% of mothers and 0.09% of fathers), parents who belong to two households at the same time (0% of mothers and 0.07% of fathers) and finally mothers for whom the difference between their youngest and second-youngest child is larger than 10 years (2.3%). We also drop observations for fathers where couples divorce or separate during the observation period (7.5%) to ensure that the partner remains the same. Finally, in the main analyses we exclude people who receive income from self-employment at any point in the observation period, because we do not observe their hours worked outside of hours spent on regular employment (10.6% of the mothers and 19.5% of the fathers).

Table 4.2 shows the summary statistics for mothers (Columns 1 and 2) and fathers (Columns 3 and 4) in the treatment and the control group weighted by matching weights. Mothers in our sample are around 37 years old in the treatment group, and as expected, somewhat younger in the control group. The treatment and the control group have a similar share of highly educated mothers and ethnicity is also similar. They have slightly more than 2 children on average. In the treatment group around 68% is married or in a civil union, while in the younger control group around 65% is married. The age at which they gave birth to their first-born is very similar, as is their hourly wage. Fathers are somewhat older than mothers, close to 40 years in the treatment group and 38 in the control group. Ethnicity and number of children – slightly more than 2 on average – are very similar. Education level is also very similar. Close to 80% in the treatment group is married or is in a civil union, while this holds for around 75% in the control group. The hourly wage is somewhat higher for the older fathers in the treatment group.

[18]Unfortunately same-sex couples cannot easily be identified in our data on household composition, because it is often unclear whether two people of the same gender are a couple or just living together in the same household. Hence we only include couples where the gender differs between the partners.

Table 4.2: Descriptives of demographics for mothers and fathers in treatment and control group weighted by matching weights.

	Mothers		Fathers	
	(1)	(2)	(3)	(4)
	Treatment group	Control group	Treatment group	Control group
Age	36.80	34.65	39.97	37.90
	(4.49)	(4.58)	(4.90)	(5.06)
High educated	0.41	0.42	0.47	0.47
	(0.49)	(0.49)	(0.50)	(0.50)
Native	0.75	0.74	0.80	0.78
	(0.43)	(0.44)	(0.40)	(0.41)
Foreign born	0.17	0.18	0.14	0.15
	(0.38)	(0.39)	(0.35)	(0.36)
Native born from foreign-born parent	0.07	0.08	0.06	0.07
	(0.26)	(0.27)	(0.24)	(0.25)
Number of children in the household	2.15	2.17	2.18	2.16
	(0.88)	(0.90)	(0.84)	(0.86)
Not married	0.20	0.24	0.20	0.25
	(0.40)	(0.43)	(0.40)	(0.43)
Married	0.68	0.66	0.80	0.75
	(0.47)	(0.47)	(0.40)	(0.43)
Single parent	0.12	0.10		
	(0.33)	(0.30)		
Age at first born	28.98	28.80	32.23	32.21
	(4.73)	(4.76)	(5.11)	(5.20)
Hourly wage pre-treatment	17.01	16.62	21.65	20.66
	(10.97)	(10.39)	(17.29)	(16.59)
Hourly wage year 1 post-treatment	17.13	16.80	21.99	21.08
	(10.77)	(9.79)	(21.41)	(16.45)
Hourly wage year 2 post-treatment	17.30	17.04	22.36	21.48
	(11.73)	(10.44)	(54.72)	(15.80)
Observations	21,473,784	21,809,088	15,364,404	16,056,612
No. of individuals	596,494	605,808	426,789	446,017

Notes: The table reports means and standard deviations in parentheses. We observe all variables for the full sample, except for education, which is only observed for about two-thirds of the sample and is biased towards higher-educated.

Figure 4.3 reports descriptive statistics on the average number of hours mothers and fathers worked and their employment status, both in the treatment and the control group, weighted by matching weights. On average, mothers work almost 16 hours per week (including those who do not work as zero). In addition, around 68% of the mothers in our sample work in a given month pre-treatment. Finally, employed mothers work around 23 hours on average per week. This suggests that there is substantial room for increases along both the intensive and the extensive margin. It is clear that the matching procedure succeeded in creating very similar groups in the pre-treatment period. This should alleviate some concerns about the common trend assumption.

Around 90% of fathers are employed. Those who are employed, work close to 38 hours per week on average. The average number of hours fathers worked declined somewhat during the three years they are in the sample. This decline is similar for the treatment and control group and seems to be driven by a decline in employment status. This could be due to the Great Recession in 2009 – 2010, and the second dip the Netherlands experienced in 2011 – 2013.

4.5 Results

In Section 4.3, we calculated that the treatment mothers experience when their youngest child starts going to school is a time shock of on average 13.3 hours per week and a monetary shock to the household of on average 61 euros per month. Similarly, fathers and stepfathers on average experience a time shock of 3.7 hours and a monetary shock to the household of 70 euros. Furthermore, Section 4.4 showed based on descriptive statistics that employment among mothers seems to increase after treatment has taken place, while the treatment hardly seems to affect men. In this section, we show the results of the difference-in-differences analyses, which confirm the descriptive results.

4.5.1 Main results

Figure 4.4 reports estimates of the effect of the youngest child going to school (at $t = 0$) on maternal and paternal employment. We present estimates for each month from the year before until two years after the child goes to school. The estimates are relative to the situation eleven months before the child goes to school.

Figure 4.4a reports the estimated differences between individuals in the treatment and the control groups on total working hours, including those who do not work as zero. The precisely estimated average effect is a small increase in weekly working hours after the youngest child goes to school of close to 0.5 hours per week after two years. This is a 3% increase relative to the pre-treatment mean of around 15 hours per week of the mothers in the treatment group. Over the entire post-treatment period the average effect is 0.3 hours, or 1.9% relative to the mean.

About 32% of the mothers in our sample does not work before their child goes to school. It is possible that the time that they no longer need to spend on childcare, induces some of them to enter the labor force. Figure 4.4c reports effects on the probability to work. We find an increase of about 1.5 percentage points relative to

Figure 4.3: Descriptive statistics on employment for mothers and for fathers

(a) Mothers' weekly working hours (including zeros)

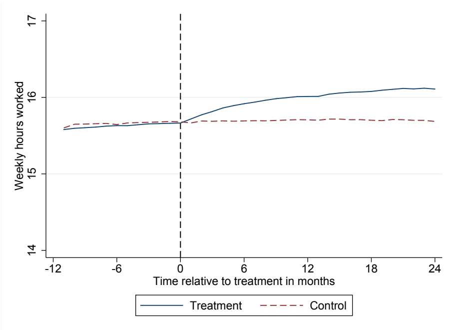

(b) Fathers' weekly working hours (including zeros)

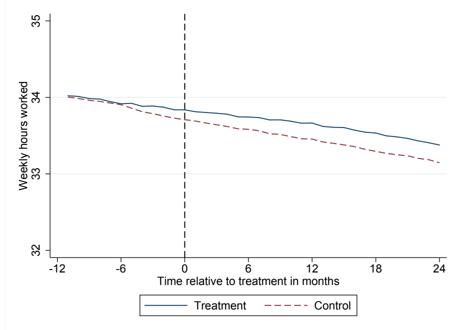

(c) Probability for mothers to be employed

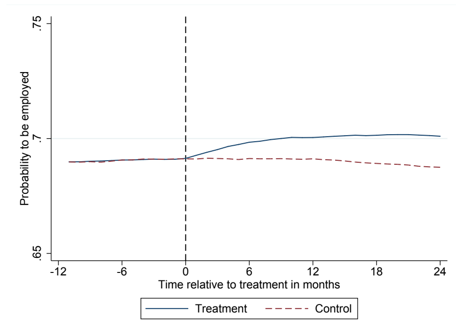

(d) Probability for fathers to be employed

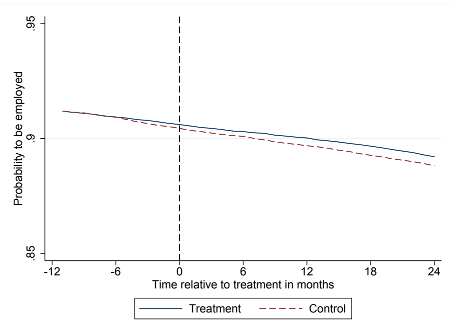

(e) Mothers' weekly working hours if employed

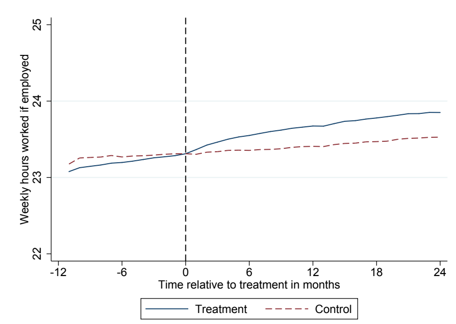

(f) Fathers' weekly working hours if employed

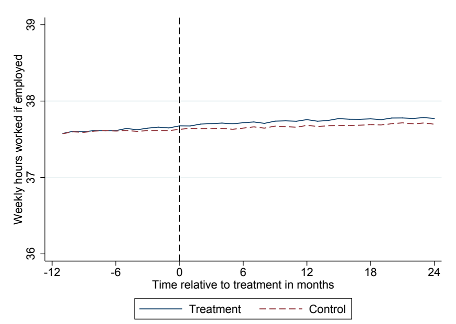

Notes: Own calculations based on register data from Statistics Netherlands. Treatment at $t = 0$. In the treatment group, this is when the youngest child turns four years old and in the control group, this is when the youngest child turns 2 years old.

the control group after two years in the chance that a mother works. This is a 2.2% increase relative to the pre-school mean of 69% of mothers who work. The average treatment effect is 0.9 percentage points, or 1.3% relative to the mean.

Finally, Figure 4.4e reports effects on weekly working hours for those in employment throughout the entire observation period. We find that those in employment work about 0.4 additional hours per week after their child goes to school. The figures provide evidence that our matching procedure succeeded in making our treatment and control groups comparable. We find no evidence of changes in labor supply in the pre-treatment period. This implies that there does not seem to be an anticipation effect.

We can now calculate the contribution of the changes at the intensive and extensive margin to the total effect. The contribution of the effect at the intensive margin is equal to the share who work multiplied by the effect size $0.68 \times 0.3 = 0.2$. The total effect is 0.3, which means that the effect measured in hours worked per week at the extensive margin must be 0.1. This means that changes along the intensive margin contribute about two-thirds to the total effect, while changes along the extensive margin contribute about one-thirds. Furthermore, given that we find an increase along the extensive of 0.9 percentage point, this means that mothers who start working when their youngest child goes to school on average work around 11 hours (0.1/0.009). This is about half the number of hours worked by those who already work before their child goes to school.

Next, we consider how the youngest child in a household going to school affects the partners of these mothers, if they have one during the observation period. Figure 4.4 shows that on average fathers and stepfathers respond much less than mothers do to their child going to school. We find that weekly working hours increase by about 0.15 hours on average, or a 0.4% increase relative to the mean of 34 hours before treatment. The probability to work also increases slightly by about 0.5 percentage points, or a 0.6% increase relative to the mean of 91% before treatment. Finally, the effect on hours worked conditional on employment is very close to zero for fathers. For fathers the effect along the intensive margin ($0.91 \times 0.05 = 0.05$) contributes about one-thirds to the total effect, while the effect along the extensive margin contributes about two-thirds. This means that fathers who start working after their youngest child goes to school, work about 20 hours on average ($0.1/0.005 = 20$). However, while statistically significant, the effects are very small. Also note that for fathers we observe a small pre-trend, suggesting that other factors could also play a role here.

Figure 4.4: Main estimates for mothers and fathers

(a) Effect on weekly working hours for mothers (including zeros)

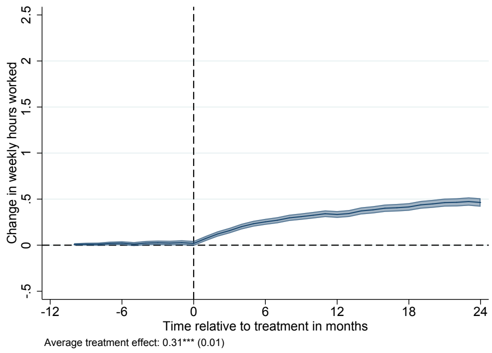

(b) Effect on weekly working hours for fathers (including zeros)

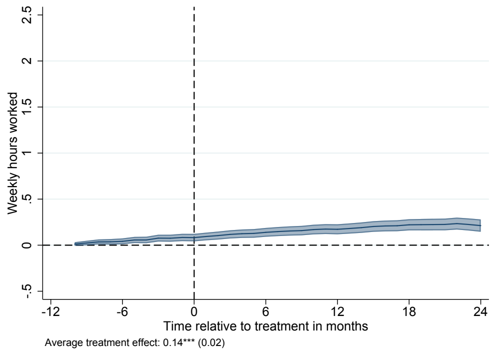

(c) Effect on probability to work for mothers

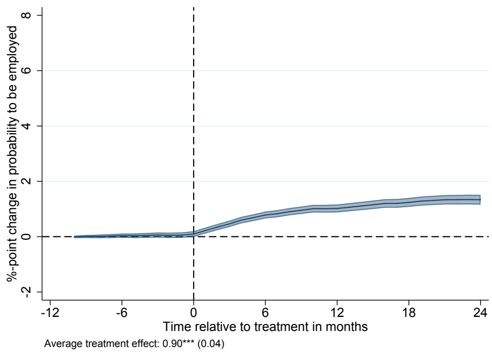

(d) Effect on probability to work for fathers

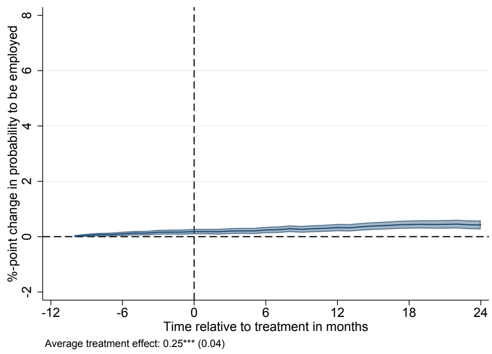

(e) Effect on weekly working hours if employed for mothers

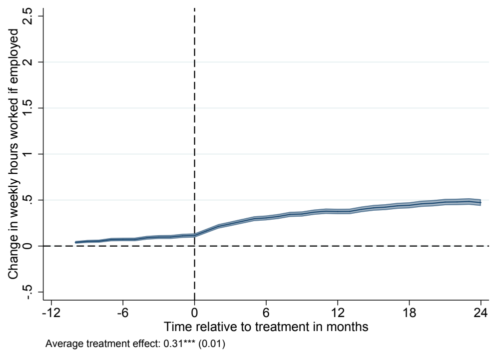

(f) Effect on weekly working hours if employed for fathers

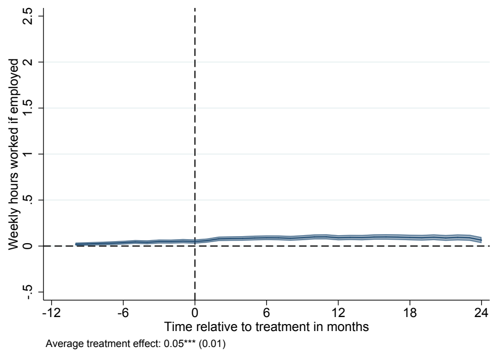

Notes: Own calculations based on register data from Statistics Netherlands. Treatment at $t = 0$. For the treatment group this is when the youngest child turns four years old. For the control group this is when the youngest child turns 2 years old. The lighter band around the estimate represents the 99% confidence interval. Reported average treatment effect is based on Equation 4.8. Cluster-robust standard errors clustered by individual in parentheses, *** p<0.001, ** p<0.01, * p<0.05. The coefficients for probability to work are multiplied by 100 so that they can be interpreted as percentage point changes.

Overall, these results show that the treatment leads to a small increase in time spent at work for mothers and an even smaller increase for fathers. The increase for mothers is both along the intensive and the extensive margin. At the same time, although the effects are statistically significant, they are economically quite small. They are also small compared to the magnitude of the change in time mothers experience when their youngest child starts going to school. Fathers generally experience a substantially smaller shock, particularly in terms of time, when their youngest child starts going to school than mothers do. This may contribute to the finding that mothers respond more than fathers when the youngest child in a household starts going to school. We look into this in the next section.

4.5.2 Heterogeneity

On average the treatment involves an increase of 13.3 hours of additional time and 61 euros of income per month for mothers and an increase of 3.7 hours and 70 euros for fathers. There is however important heterogeneity in the treatment as was shown in Table 4.1. Parents who work more hours gain less additional time than parents who work very few hours per week. Additionally, households in which the parents earn more, gain more money when their youngest child goes to school than household with a lower household income because childcare subsidies are income dependent.

Theoretically, we would expect larger positive responses in labor supply for those with the largest time gains and the smallest financial gains. However, our heterogeneity analyses come to different conclusions. Differences in initial preferences for labor supply are probably at the heart of these results.

Table 4.3 shows the differences in the estimated effects on the hours worked (including those who do not work) for mothers (Column 1) and fathers (Column 3) who work different numbers of hours (Panel A) and who earn different wages per hour (Panel B) in the year before treatment. Column 2 (Column 4) reports the mean number of hours mothers (fathers) in the treatment group worked in the twelve months before their youngest child started going to school. Panel A shows that –surprisingly– the increase in hours worked is largest for mothers who already worked more than 34 hours before treatment. The increase in hours worked is smallest among those who did not work before the treatment. So although the average number of hours mothers worked increases significantly amongst all groups, mothers who obtain more additional time (because they work fewer hours) when their youngest child starts going to school in fact increase their working hours *less* than mothers who only obtain very little additional time (because they already worked (nearly) fulltime before the treatment).

On average, fathers experience a substantially smaller shock in terms of time when their youngest child starts going to school. Column 3 of Table 4.3 illustrates that fathers who do experience a larger shock in terms of time (because they work few hours prior to the treatment), also do not significantly increase their working hours.These findings suggest that there are reasons other than caring for the youngest child to explain why these fathers do not work, such as disability or study. These results point out that the differences between the estimated average treatment effects between mothers and fathers are not simply entirely due to differences in the magnitude of the shock they experience when their youngest child goes to school. Even within groups in which the number of hours parents work is relatively similar, mothers increase their working hours more or at least as much as fathers do.

Panel B reports the heterogeneity by the hourly wage parents earn. In Section 4.4, we formulated the hypothesis that for parents who earn higher wages, the effects are larger. Clearly wages are only defined for parents who are working. For mothers, we see that women who earn higher wages on average work more hours, while this does not seem to be the case for fathers. We find that all groups increase their hours worked, and that for both mothers and fathers, the increase is indeed largest among those who earn higher hourly wages.

Appendix Table A2 additionally reports how the effects differ by demographic characteristics. Contrary to what is commonly found in the literature, we see in Panel A that single mothers actually decrease the number of hours they work when their youngest child starts going to school. This might be explained by differences in the examined time periods: most papers on the effect of compulsary schooling on maternal labor supply use data from the 1980s and the 1990s and in those days the composition and preferences of the female labor force were distinctly different. At the same time, the problem might also be that our control group is not suitable for single mothers: perhaps single mothers with very young children are inherently different from single mothers with slightly older children.

Panel C shows that mothers with one child show a small decline in hours worked and the probability to work, while mothers with more children increase both their hours and participation in the labor market. One explanation could be that mothers with more children wait until their youngest is off to school, and only then enter the labor market, while mothers with only one child were already working more. Indeed, the data show that mothers with one child work 17.7 hours on average, while

Table 4.3: Heterogeneity by hours worked and wage quartile.

	Mothers		Fathers	
	(1)	(2)	(3)	(4)
	Estimated effect	Avg hours worked	Estimated effect	Avg hours worked
Panel A. Heterogeneity by hours worked				
Not working	0.22***	0.0	−0.01	0.0
	(0.02)		(0.09)	
1-20 hours	0.35***	12.9	0.12	10.8
	(0.02)		(0.14)	
20-33 hours	0.31***	25.5	0.31***	29.9
	(0.01)		(0.05)	
34+ hours	0.50***	37.9	0.13***	38.8
	(0.06)		(0.01)	
Observations	43,282,872		31,421,016	
Panel B. Heterogeneity wage quartile				
Quartile 1	0.22***	19.1	0.06	36.1
	(0.03)		(0.03)	
Quartile 2	0.36***	21.3	0.08***	37.1
	(0.02)		(0.02)	
Quartile 3	0.44***	22.4	0.10***	36.8
	(0.02)		(0.02)	
Quartile 4	0.45***	25.5	0.24***	37.0
	(0.02)		(0.03)	
Observations	29,785,932		28,631,088	
Panel C. Full sample				
Average treatment effect	0.31***	15.6	0.14***	33.9
	(0.01)		(0.02)	
Observations	43,282,872		31,421,016	

Notes: The table reports the total estimated effect for each group from a regression with interactions between the treatment indicator and indicators for each groups. Cluster-robust standard errors clustered by individual in parentheses, *** $p<0.001$, ** $p<0.01$, * $p<0.05$. All regressions include individual fixed effects and calendar year-month fixed effects.

mothers with more than two children work 16.9 hours on average and mothers with three or more children work 11.6 hours on average.

4.5.3 Robustness analyses

In the main analyses, we did not include self-employed workers in our sample, because we do not observe the number of hours they worked. Because this restriction reduces the sample by a sizable 10.6% for mothers and 19.5% for fathers, it is important to discuss whether this restriction affects our results.

To check this, we estimate the hours self-employed mothers (fathers) worked in a given year by dividing the profits they reported by the average hourly wage of all mothers (or fathers) in our sample in a given calendar year.[19] Clearly using the average hourly wage to calculate the number of hours self-employed parents work yields only a rough estimation of the actual number of hours a self-employed parent works. Still, the estimates for the probability to work are reliable. Other than excluding the self-employed parents, we apply the same sample selections as in the main analysis.[20] Descriptives are reported in Table A3 in the Appendix.

Table 4.4 illustrates that including self-employed in the sample does not change the main conclusions. Panel C reproduces the original estimates from the main analyses and Panel A shows how the effects change when self-employed people are included in the sample. For mothers, we find effects very similar to the main estimates. Including self-employed, mothers work about 1.8% more hours on average in the two years after their youngest child starts going to school, compared to 2% for the main estimates. The estimates for hours worked for those who are employed and for the extensive margin are also remarkably similar. For fathers, we find an average increase in total hours worked of 0.3% relative to the mean, compared to an increase of 0.4% for the main estimates. We find no statistically significant effect for those in work, which means that the full increase is driven by the extensive margin. The results on the extensive margin are very similar to the main estimates. This suggests that the substantial sample selection of self-employed does not bias our results.

4.5.3.1 An alternative control group

To further check the robustness of our results, we also consider an alternative control group. In our main analyses, our control group consists of parents whose youngest child is between one and four years old. Their youngest child does not go to school in the observation period. An alternative approach is to use parents whose second-youngest child is between three and six years old as a control group.[21] In line with

[19] We therefore assume that these self-employed parents are working in each month throughout the calendar year if they have income from self-employment in that year. If self-employed workers are also employed at some point in time, we simply sum the hours worked as self-employed and in their employment.

[20] We again drop people who work more than 100 hours per week in either a job, or as self-employed or as a combination of both. For mothers this is 0.7% of the sample and for fathers it is 1.5%. We then run the same matching procedure as described in the main text. We end up with 1,341,925 matched mothers and 1,127,896 matched fathers. This is 11.6% (29.2%) more than the main sample. 66% (91%) of the mothers (fathers) in this sample work in a given month before treatment, and they work 16.3 (33.7) hours on average.

[21] To ensures that parents in this group do not experience a time shock because their youngest child starts going to school while they are in this control group, the youngest child should be two or three years younger and the child remains the second-youngest child in the observation period.

the parents in the control group in our main analyses, parents in this alternative control group do not experience a time shock during the observation period, because they need to take care of their youngest child who does not go to school yet. Parents in this alternative control group, however, do experience a small monetary shock, because they no longer need to pay for child care in the hours during which their second-youngest child now starts going to school. In the Netherlands, government subsidies are substantially higher for the child in the household that uses most hours of child care (generally the youngest child) than for other children. As a result, the monetary shock households experience when their second-youngest child starts going to school are substantially lower. Appendix Table A1 shows the distribution of the magnitude of the shock fathers and mothers in this alternative control group experience.

We apply the same selections as in the baseline estimates. We end up with a sample of 596,494 (426,789) treated mothers (fathers) and 152,947 (122,012) control mothers (fathers). Descriptives are reported in Table A4 in the Appendix.

Panel B in Table 4.4 reports the estimation results. For mothers we find somewhat smaller effects on average, but larger effects along the intensive margin. We find no effect along the extensive margin. For fathers we actually find small negative effects on hours worked overall. These appear to be driven by the extensive margin. However, while statistically significant, the results are, just as the baseline estimates, economically insignificant: we find a decline in hours worked overall of 0.5% for fathers. Hours worked while employed increases by 0.05 hours, or 0.1%, just as in the baseline estimates. We also find a small decline in probability to work of about 0.6%.

4.6 Conclusion and discussion

This paper shows that when the youngest child in a family starts going to school, mothers in the Netherlands only increase their labor supply marginally. When the youngest child starts attending school, mothers experience an increase in free time of more than thirteen hours a week, which is expected to increase their labor supply. Theoretically, this effect may be mitigated by the income effect that households experience simultaneously: when the youngest child starts attending school, parents save on average sixty to seventy euros on the costs of formal child care. Nonetheless, even though such income effects are expected to decrease parental labor supply, the

Table 4.4: Treatment effect estimates on hours per week, hours worked if employed and the probability to work when including self-employed and relative to a control group consisting of parents of children with a second-youngest child aged 3–6.

	Mothers			Fathers		
	(1)	(2)	(3)	(4)	(5)	(6)
	Hours worked	Hours worked if working	Probability to work	Hours worked	Hours worked if working	Probability to work
Panel A. Including self-employed						
Average treatment effect	0.29^{***}	0.28^{***}	0.87^{***}	0.10^{***}	0.02^{*}	0.24^{***}
	(0.01)	(0.01)	(0.04)	(0.02)	(0.01)	(0.04)
Observations	48,309,300	26,732,808	48,309,300	40,604,256	31,022,640	40,604,256
Mean dep. variable	16.2	23.7	0.70	33.7	37.5	0.91
Panel B. Relative to control group with second-youngest child 3-6						
Average treatment effect	0.23^{***}	0.45^{***}	0.09	-0.15^{***}	0.08^{***}	-0.54^{***}
	(0.02)	(0.01)	(0.07)	(0.02)	(0.01)	(0.06)
Observations	26,979,876	15,414,624	26,979,876	19,756,836	15,759,468	19,756,836
Mean dep. variable	15.6	23.2	0.69	33.9	37.6	0.91
Panel C. Main results						
Average treatment effect	0.31^{***}	0.31^{***}	0.90^{***}	0.14^{***}	0.05^{***}	0.25^{***}
	(0.01)	(0.01)	(0.04)	(0.02)	(0.01)	(0.04)
Observations	43,282,872	23,990,148	43,282,872	31,421,016	24,695,748	31,421,016
Mean dep. variable	15.6	23.2	0.69	33.9	37.6	0.91

Notes: Cluster-robust standard errors clustered by individual in parentheses, *** $p<0.001$, ** $p<0.01$, * $p<0.05$. All regressions include individual fixed effects and calendar year-month fixed effects. The treatment group consists of mothers with a youngest child 3–6 years old. The control group in Panels A and C consists of mothers with a youngest child between 1 and 4 years old. In Panel B the control group consists of parents with a second-youngest child 3–6 years old. The coefficients for probability to work are multiplied by 100 so that they can be interpreted as percentage point changes.

effects are likely to be small in this case, because the monetary shock is relatively small.

Compared to mothers whose youngest child turned two and who therefore do not experience any shock in the observation period, mothers increase their labor supply by on average approximately 0.5 hours a week after two years. This is around a 3% increase relative to the mean before their child goes to school. For their partners we find an even smaller increase of about 0.15 hours after two years, or 0.4% relative to the mean before their child goes to school. These findings are generally robust to alternative specifications. However, employing an alternative control group of parents with second-youngest children leads to small negative estimates for partners and somewhat smaller positive effects for mothers.

Our heterogeneity analyses shows surprising results: the response is larger among mothers and fathers who are already working longer hours and who already earn higher wages. Theoretically, one would expect smaller increases in working hours in these groups as the relative gain in time and opportunity costs of child care are smaller. However, differences in initial preferences for labor supply over leisure are probably driving these results.

Utility or consumption smoothing can perhaps be an explanation for the relatively limited effects that we find in this paper. Clearly, the moment a child turns four is known in advance. Assuming that parents wish to keep their marginal utility of consumption and leisure as similar as possible over time, only small changes to the number of hours worked can be expected. An alternative explanation is that parents in fact consider the gross monetary gain - the difference in the cost of child care without considering the corresponding decline in child care subsidies - they obtain when their youngest child starts going to school, rather than the net monetary gain that we calculate in this paper. In the Netherlands parents receive child care subsidies from the tax authorities and often on a different date than when they pay for child care. The gross monetary gain is substantially larger than the net monetary gain, especially for parents with a relatively low household income, and this leads to a larger perceived increase in income than is actually the case. In addition, the literature shows that successive cohorts of parents respond less and less strongly to incentives aimed at increasing their labor participation. Finally, it could be that school times are inconvenient, making it difficult to increase hours worked during school hours.

A Appendix

A.1 Description of maternal budget constraint

This appendix explains in more detail how the time mothers spend on work, taking care of their children and leisure changes when the youngest child starts going to school, i.e. when $t_s > 0$. Figure 4.2 shows the budget constraint mothers face when at least one child in a household does not go to school and the budget constraint mothers face when the youngest child in a household starts going to school. In this figure the time fathers spend on labor l_f and child care t_f is held constant. Mothers decide how to allocate their time between leisure z_m, taking care of their children t_m and working from which they earn an income that enables them to consume X. The blue solid planes represent the situation when at least one child does not go to school yet and the green dashed planes represent the situation after the youngest child has started going to school when parents no longer have to arrange child care during school hours.

Figure 4.2a demonstrates the tradeoff between the time mothers take care of their children t_m and the time they spend on leisure z_m, Figure 4.2b depicts the tradeoff between the mother's leisure z_m and the household's consumption X and Figure 4.2c shows the relation between the time mothers spend on domestic child care and household consumption.

In this section we explain the budget constraint by zooming in on four situations that correspond to four points on this budget constraint. A first situation of interest is when mothers work fulltime and therefore do not enjoy any free time and do not provide domestic child care (1: $l_m = 1, z_m = 0, t_m = 0$). A second situation is when mothers spend all their time on leisure ($l_m = 0, z_m = 1, t_m = 0$). At the third and fourth situation we discuss, the mother spends as much time as she can on domestic child care. When a child goes to school or when the father (whose time use is held constant in Figure Figure 4.2) also provides child care, mothers cannot spend all their time on domestic child care. The third situation therefore entails that women spend as much time as possible on providing domestic child care and spend the remainder of their time working (3A: $l_m = t_s + t_f, z_m = 0, t_m = 1 - t_s - t_f$). Finally in the fourth situation the mother again maximizes the time she spends on providing domestic child care and spends the remainder of her time on leisure (3B: $l_m = 0, z_m = t_s + t_f, t_m = 1 - t_s - t_f$).

First, when mothers work fulltime and do not enjoy any free time and do not provide domestic child care (1: $l_m = 1, z_m = 0, t_m = 0$), the household pays for child care while the parents work and consumes the earnings of the father and the mother. Household consumption therefore equals $X = w_f^* l_f + w_m^* - q^*(1 - t_s - t_f)$, where the asterisks denote real wages and real prices of child care, $w_i^* = w_i/p$ and $q^* = q/p$. In Figure 4.2a this point is at the origin: mothers do not provide child care nor do they enjoy leisure. In Figure 4.2b and in Figure 4.2c this situation is at the top-left part of the budget constraint. As a result, there is a linear upward shift of the budget constraint from the blue plane to the green dashed plane at this point. When the youngest child in a household starts going to school, household consumption increases by $q^* t_s$ and depending on the shape of the utility function this causes mothers to increase their leisure and thereby decrease their labor supply.

A second situation of interest is when mothers spend all their time on leisure. In this situation, women do not work but children attend external child care outside of possible school hours and when their father does not take care of them. Household consumption at this point is therefore $X = w_f^* l_f - q^*(1 - t_s - t_f)$. In line with the previous situation, there is an upward linear shift of the budget constraint when the youngest child starts going to school and parents only experience an income effect.

Third, when mothers spend as much time as they can on providing domestic child care, they additionally experience a time effect. In situation 3A, mothers maximize the time they spend on child care and work the remainder of the time, $l_m = t_s + t_f, z_m = 0, t_m = 1 - t_s - t_f$. In this situation, children do not attend external child care. Household consumption then equals $X = w_f^* l_f + w_m^*(t_s + t_f)$. When the youngest child starts attending school, the time these mothers spend on domestic child care decreases from $1 - t_f$ to $1 - t_s - t_f$ and her labor supply increases correspondingly.

Finally, in situation 3B mothers spend as much time as possible on providing domestic child care and spend the remainder of the time on leisure, $l_m = 0, z_m = t_s + t_f, t_m = 1 - t_s - t_f$. Household consumption in this situation is equal to $X = w_f^* l_f$ as children do not attend external child care and mothers do not work. When the youngest child starts going to school, these mothers spend less time on domestic child care and enjoy more leisure.

In sum, when the youngest child starts going to school, the budget constraint changes in two ways: on the one hand there is an upward linear shift because mothers whose children attend external child care experience an income effect as their child care expenses decrease. On the other hand, the maximum amount of time mothers can spend on domestic child care decreases. The exact optimal situation in the situation before and after the youngest child starts going to school depends on the marginal utility she obtains from those activities, and thereby on the shape of the utility function.

A.2 Additional results

Table A1: Magnitude of the shock in terms of money and in terms of time

Panel A: Shock in terms of time and money that mothers experience when their second-youngest child goes to school

	(1) % of mothers	(2) Hours saved per week	(3) Money saved per month (per household)
Does not work	19.2	0	3
Works less than 20 hours a week	27.3	0	11
Works 20-34 hours a week	48.6	0	28
Works more than 34 hours a week	4.9	0	48
Total	100.0	0	19

Panel B: Shock in terms of time and money that fathers experience when their second-youngest child goes to school

	% of fathers	Hours saved per week	Money saved per month (per household)
Does not work	4.8	0	9
Works less than 20 hours a week	3.4	0	14
Works 20-34 hours a week	18.1	0	23
Works more than 34 hours a week	73.6	0	21
Total	100.0	0	21

Notes: See main text for calculation of the monetary and time shocks. The percentages listed in Column (1) refer to the share of mothers (fathers) in our sample that works a specific number of hours. The amounts of money households save when their second-youngest child starts going to school (listed in Column (3)) differs for the mothers (Panel A) and fathers (Panel B) in our sample, because households with single mothers *are* included in the sample for mothers, but not in the sample of fathers.

Table A2: Heterogeneity by marital status, ethnicity and number of children.

	Mothers		Fathers	
	(1)	(2)	(3)	(4)
	Estimated effect	Avg hours worked	Estimated effect	Avg hours worked
Panel A. Heterogeneity by marital status				
Dual household	0.37***	15.9		
	(0.01)	(0.01)		
Single household	−0.29***	13.2		
	(0.05)	(0.04)		
Panel B. Heterogeneity by ethnicity				
Native	0.40***	16.8	0.07***	35.1
	(0.01)		(0.02)	
Foreign	0.05*	12.1	0.30***	29.6
	(0.02)		(0.04)	
Panel C. Heterogeneity by number of children				
1 child	−0.17***	17.7	0.18***	33.2
	(0.03)		(0.04)	
2 children	0.45***	16.9	0.14***	34.7
	(0.01)		(0.02)	
3 or more children	0.41***	11.6	0.10***	32.7
	(0.02)		(0.03)	
Panel D. Full sample				
Average treatment effect	0.31***	15.6	0.14***	33.9
	(0.01)		(0.02)	
Observations	43,282,872		31,421,016	

Notes: The table reports the total estimated effect for each group from a regression with interactions between the treatment indicator and indicators for each groups. Cluster-robust standard errors clustered by individual in parentheses, *** $p<0.001$, ** $p<0.01$, * $p<0.05$. All regressions include individual fixed effects and calendar year-month fixed effects.

Table A3: Descriptives of demographics for mothers and fathers in treatment and control group weighted by matching weights for the sample including self-employed.

	Mothers		Fathers	
	(1)	(2)	(3)	(4)
	Treatment group	Control group	Treatment group	Control group
Hours worked per week pre-treatment	16.26	16.20	33.58	33.46
	(13.58)	(13.42)	(14.02)	(13.93)
Fraction working pre-treatment	0.70	0.70	0.91	0.91
	(0.46)	(0.46)	(0.29)	(0.29)
Hourly wage pre-treatment	17.15	16.77	22.43	21.33
	(11.61)	(11.54)	(25.53)	(19.80)
Hourly wage year 1 post-treatment	17.25	16.95	22.69	21.66
	(10.99)	(11.01)	(25.87)	(18.95)
Hourly wage year 2 post-treatment	17.42	17.17	22.98	22.02
	(12.33)	(11.06)	(53.92)	(18.87)
Self-employed at some point	0.11	0.10	0.23	0.22
	(0.31)	(0.30)	(0.42)	(0.41)
Age	36.91	34.77	40.06	38.03
	(4.47)	(4.56)	(4.89)	(5.06)
Fraction high educated	0.43	0.44	0.47	0.47
	(0.49)	(0.50)	(0.50)	(0.50)
Native	0.76	0.74	0.80	0.79
	(0.43)	(0.44)	(0.40)	(0.41)
Foreign born	0.17	0.18	0.13	0.14
	(0.37)	(0.38)	(0.34)	(0.35)
Native born from foreign-born parent	0.07	0.08	0.06	0.07
	(0.26)	(0.27)	(0.24)	(0.25)
Number of children	2.16	2.17	2.21	2.20
	(0.88)	(0.90)	(0.86)	(0.88)
Not married	0.20	0.25	0.21	0.26
	(0.40)	(0.43)	(0.41)	(0.44)
Married	0.68	0.66	0.79	0.74
	(0.46)	(0.47)	(0.41)	(0.44)
Single parent	0.12	0.09		
	(0.32)	(0.29)		
Age at first born	29.08	28.93	32.24	32.25
	(4.71)	(4.76)	(5.15)	(5.25)
Observations	24,117,624	24,191,676	20,004,552	20599704
No. of individuals	669,934	671,991	555,682	572,214

Notes: The table reports means and standard deviations in parentheses. We observe all variables for the full sample, except for education, which is only observed for about two-thirds of the sample and is biased towards higher-educated.

Table A4: Descriptives of demographics for mothers and fathers in treatment and control group weighted by matching weights for the sample using a control group of parents whose second-youngest child is between three and six years old.

	Mothers		Fathers	
	(1)	(2)	(3)	(4)
	Treatment group	Control group	Treatment group	Control group
Hours worked per week pre-treatment	15.67	15.71	33.89	33.90
	(12.65)	(12.57)	(12.29)	(12.11)
Fraction working pre-treatment	0.69	0.69	0.91	0.91
	(0.46)	(0.46)	(0.29)	(0.29)
Hourly wage pre-treatment	17.01	17.52	21.71	21.47
	(10.97)	(11.81)	(17.14)	(16.52)
Hourly wage year 1 post-treatment	17.13	17.73	21.99	21.87
	(10.77)	(10.32)	(21.41)	(18.71)
Hourly wage year 2 post-treatment	17.30	18.00	22.35	22.32
	(11.73)	(12.10)	(56.77)	(15.35)
Age	36.80	34.17	39.97	37.02
	(4.49)	(4.13)	(4.90)	(4.58)
High educated	0.41	0.52	0.47	0.55
	(0.49)	(0.50)	(0.50)	(0.50)
Native	0.75	0.80	0.80	0.84
	(0.43)	(0.40)	(0.40)	(0.37)
Foreign born	0.17	0.13	0.14	0.10
	(0.38)	(0.34)	(0.35)	(0.30)
Native born from foreign-born parent	0.07	0.07	0.06	0.06
	(0.26)	(0.25)	(0.24)	(0.24)
Number of children	2.15	2.38	2.18	2.31
	(0.88)	(0.75)	(0.84)	(0.66)
Not married	0.20	0.24	0.20	0.25
	(0.40)	(0.43)	(0.40)	(0.43)
Married	0.68	0.71	0.80	0.75
	(0.47)	(0.45)	(0.40)	(0.43)
Single parent	0.12	0.05	0.00	0.00
	(0.33)	(0.21)	(0.00)	(0.00)
Age at first born	28.98	28.97	32.23	32.06
	(4.73)	(4.22)	(5.11)	(4.53)
Observations	21,473,784	5,506,092	15,364,404	43,924,32
No. of individuals	596,494	152,947	426,789	122,012

Notes: The table reports means and standard deviations in parentheses. We observe all variables for the full sample, except for education, which is only observed for about two-thirds of the sample and is biased towards higher-educated.

CHAPTER 5

Using Tax Deductions to Promote Lifelong Learning: Real and Shifting Responses*

5.1 Introduction

Lifelong learning is high on the policy agenda. Societal and technological changes increase the need to invest in lifelong learning. For example, effective retirement ages in developed economies have risen dramatically over the past decade.[1] Also, technological change and globalization seem to reduce the lifespans of sectors, firms and products (Goos et al., 2014; Autor, 2015). As a result, individuals are more likely to switch jobs and careers during their working life, and are more likely to switch tasks within a given job. In the face of these changes, maintaining and investing in human capital during the working life becomes increasingly important. At the same time, policymakers worry that individuals and/or their employers underinvest in lifelong learning, due to e.g. hold-up problems (Malcomson, 1997, 1999).[2] Although it is difficult to determine empirically whether there is underinvestment in lifelong learning in general, policymakers seem particularly worried about certain subgroups

*This chapter is joint work with Egbert Jongen and Karen van der Wiel. It is based on Van den Berge et al. (2017).

[1]The Netherlands is no exception and current 30-year olds are expected to retire beyond their 70th birthday.

[2]Though studies have also identified factors that may mitigate this hold-up problem, like reciprocity and smart contract designs (Leuven et al., 2005; Hoffman and Burks, 2013).

of the population that have a distaste for formal learning, such as lower educated individuals (see e.g. Eurostat, 2016) and workers in sectors that seem particularly 'at risk' due to technological change and globalization. Policymakers therefore try to mitigate potential underinvestment in lifelong learning, by providing financial support to employees and their employers that undertake lifelong learning, regulating and funding post-initial education and training, informing employees and their employers about the possibilities for lifelong learning and scrutinizing labor market regulations for adverse side effects on lifelong learning. Recently, a literature has emerged that investigates the effectiveness of different policy measures. However, so far only direct financial support measures have been investigated systematically and even then the empirical evidence on the effectiveness of this type of policy remains scarce. On the prospects for tax incentives to stimulate lifelong learning we know very little.

In this paper we study whether a tax deduction for lifelong learning can stimulate investment in lifelong learning. Specifically, we consider the effects of a tax deduction in the Netherlands, where individuals can deduct their expenditures on post-initial work-related training and education from their pre-tax personal income. Jumps in marginal tax rates provide exogenous variation in the financial incentives to undertake lifelong learning. We study the effect of this exogenous variation on the probability of filing lifelong learning expenditures and on the amount of lifelong learning expenditures filed, for different subgroups and at different points in the income distribution.

We employ a regression kink and a regression discontinuity design to estimate the impact of the tax deduction on lifelong learning expenditures. The Dutch income tax system features two discontinuous jumps in the statutory marginal tax rate. Moving from the left to the right of the discontinuity, the upward jump in the marginal tax rate implies a lower effective cost for lifelong learning to the right of the discontinuity. We prefer the regression kink design, which we can apply to singles, as the necessary conditions are met for this group. For couples, however, we observe bunching at the kink, which we address by estimating a so-called donut regression discontinuity design. We use a high quality administrative dataset of tax returns on the universe of Dutch taxpayers for the years 2006–2013. This dataset provides information on all relevant earnings activities of the Dutch population, and also contains all the information on tax deductions. A particularly unique feature of the dataset is that it contains information on the amount spent by each fiscal partner, and on the amount filed by each partner after they potentially shift part of the expenditures to the partner with the highest marginal rate.

Our main findings are as follows. First, for singles we find heterogeneous effects of the tax deduction on the probability to file lifelong learning expenditures and on the amount of lifelong learning expenditures filed.[3] The effect at the kink at a relatively low income level (approximately 18 thousand euro) is close to zero and not statistically significant. However, at the kink at a relatively high income level (approximately 55 thousand euro) the probability to file lifelong learning expenditures increases by 10%. Looking at the effects for different demographic subgroups of at the higher income kink, we find larger effects for male, native, higher educated and middle-aged singles. Second, for couples, for primary earners (partners that earn more than their partner) we find large effects on the probability to file lifelong learning expenditures and the average amount filed, at both income kinks. For secondary earners (partners that earn less than their partner) we find counterintuitive negative effects. However, we show that these results are biased, due to the shifting of the lifelong learning expenditures between partners.[4] Third, when we consider the individual lifelong learning expenditures of each partner before shifting between partners, and leave out the bins with excess mass close to the tax bracket thresholds, we find smaller effects for primary earners and the effect becomes close to zero for secondary earners.

We make a number of contributions to the literature. We contribute to the scarce literature on the causal effects of tax incentives for lifelong learning. We build on the analysis by Leuven and Oosterbeek (2012), but make substantial improvements. The authors use a sample of about 100 thousand Dutch tax returns, of which only a subsample of individuals is close to the relevant tax bracket thresholds. Our paper uses about 10 million tax returns. Furthermore, we estimate separate regressions for singles and couples, which turns out to be very important for the results, and take manipulation of the running variable into account. Furthermore, for singles we can use the more novel regression kink design approach, which seems more appropriate than the regression discontinuity used in Leuven and Oosterbeek (2012) given the kink in the financial incentive driving the result. Finally, for couples, we have the amount of lifelong learning expenditures before and after shifting between fiscal partners, whereas Leuven and Oosterbeek (2012) only had access to data on lifelong learning expenditures after shifting. Our analysis shows that ignoring shifting of

[3]Our sample of singles includes both childless singles and lone parents, what is important for our analysis is that both of these groups of 'singles' have no fiscal partner.

[4]In the elasticity-of-taxable-income literature, theoretical and empirical studies have also shown that it is important to account for tax shifting when estimating the (relevant) elasticity of the tax base with respect to the tax rate, see e.g. Kopczuk (2005), Chetty (2009), Saez et al. (2012) and Doerrenberg et al. (2017).

lifelong learning expenditures between partners leads to spurious large estimates for primary earners and spurious negative estimates for secondary earners.

The only other paper, to the best of our knowledge, to directly study the effectiveness of tax stimuli for lifelong learning, but then targeted at employers, is Leuven and Oosterbeek (2004). They find that a tax advantage for employers for training activities of their workers over the age of 40 only shifted training expenses from employees just below 40 to those just over 40, with little to no effect on overall training expenses.

Furthermore, we contribute to the general literature on the impact of financial incentives on lifelong learning. These papers typically find positive but limited effects. For example, Schwerdt et al. (2012) investigates a general voucher program in Switzerland, Hidalgo et al. (2014) look at a voucher program for specific sectors in the Netherlands and Görlitz and Tamm (2016) analyze a large co-financing instrument in Germany. In all cases, employees could pick a short training program at lower than regular costs. Training participation increased by between 13 to 20 percentage points due to these subsidies.[5] Furthermore, Schwerdt et al. (2012) also consider heterogeneous treatment effects and find that lower educated individuals seem to benefit somewhat more by participating in additional training in terms of higher wages. Other papers in this literature investigate policies in which employers receive (part of) the subsidy directly (Görlitz, 2010; Abramovsky et al., 2011; Van der Steeg and van Elk, 2015).

Our paper also relates to a relatively new literature studying the effects of tax incentives on initial education (Dynarski and Scott-Clayton, 2016). In countries with many private schools, tuition expenses can be substantial and sometimes the tax authorities are subsidizing these expenditures directly. Also, savings for future college tuition expenditures are in certain cases deductible. These tax subsidies are both meant to increase private school and college attendance, and to give income support to low- and middle income families with children. A few papers have been able to identify causal effects on higher education participation and these papers found small effects of these tax subsidies at best. Bulman and Hoxby (2015) find negligible effects on several outcomes in higher education of three tax credits for households who pay tuition and fees. Hoxby and Bulman (2016) argue that this might be due to the price inelasticity of marginal households, but that limited knowledge about the deduction and the delay in receiving the financial benefit also matter.

[5] However, no wage or employment effects were found for subsidy recipients.

The outline of the paper is as follows. Section 5.2 gives a brief description of relevant elements of the Dutch income tax system and the tax deduction for lifelong learning. Section 5.3 outlines a stylized life-cycle model that makes predictions about the relationship between the tax deduction and marginal tax rates and investments in lifelong learning, which motivates the setup of our empirical analysis. Section 5.4 discusses our empirical methodology. A description of the dataset, including descriptive statistics, is given in Section 5.5. Section 5.6 presents the main results as well as a number of robustness checks. Section 5.7 concludes. An appendix contains supplementary material.

5.2 Institutional setup

We exploit differences in marginal tax rates to identify the effect of the tax deduction on lifelong learning expenditures in the Netherlands. In this section we discuss how the tax deduction for lifelong learning works and outline the relevant characteristics of the Dutch income tax system for our sample period (2006–2013).

The tax deduction for lifelong learning is an income tax deduction for out-of-pocket expenditures on post-initial work-related training and education. The financial gain of the tax deduction is equal to the expenditures (minus a threshold) multiplied by the marginal income tax rate. The marginal income tax rate is a step-wise increasing function of individual taxable income.[6] Table 5.1 shows the marginal tax rates and tax brackets for the period 2006–2013. The difference between the tax rates in the first and second bracket is approximately 8 percentage points over the period 2006 to 2012, and drops to 5 percentage points following the increase in the marginal tax rate for the first bracket in 2013. The difference between the tax rates in the third and fourth bracket is 10 percentage points throughout the entire sample period. The beginning and end of the tax brackets have changed relatively little, they are indexed with inflation, except in 2013, when the end of the first bracket increased somewhat, while the end of the second and third brackets decreased somewhat. The change in the tax rates and tax brackets in 2013 are two reasons why we exclude 2013 from our main analyses, in addition to the changes in the deduction for lifelong learning expenditures in 2013 discussed below.

Lifelong learning expenditures are only deductible if the goal is to stimulate human capital formation and/or to improve one's labour market position. This

[6]We consider the role of income dependent tax credits and subsidies, and the relevant differences between the statutory and effective marginal tax rates, in the discussion in Section 5.6.3.

Table 5.1: Marginal tax rates and income brackets: 2006–2013

	First bracket	Second bracket	Difference	Third bracket	Fourth bracket	Difference
Bracket tax rate (%)						
2006	34.15	41.45	*7.30*	42.00	52.00	*10.00*
2007	33.65	41.40	*7.75*	42.00	52.00	*10.00*
2008	33.60	41.85	*8.25*	42.00	52.00	*10.00*
2009	33.50	42.00	*8.50*	42.00	52.00	*10.00*
2010	33.45	41.95	*8.50*	42.00	52.00	*10.00*
2011	33.00	41.95	*8.95*	42.00	52.00	*10.00*
2012	33.10	41.95	*8.85*	42.00	52.00	*10.00*
2013	37.00	42.00	*5.00*	42.00	52.00	*10.00*
Top of the tax bracket (euro)						
2006	17,046	30,631		52,228	∞	
2007	17,319	31,122		53,064	∞	
2008	17,579	31,589		53,860	∞	
2009	17,878	32,127		54,776	∞	
2010	18,218	32,738		54,367	∞	
2011	18,628	33,436		55,694	∞	
2012	18,945	33,863		56,491	∞	
2013	19,645	33,363		55,991	∞	

includes tuition fees, books, necessary clothing and depreciation on a computer when the computer is necessary for a work-related course. Living and travel expenses are excluded, as are expenditures on courses for strictly personal development, ‘hobbies’ and materials used for self study. Furthermore, untaxed benefits for lifelong learning, such as a study grant from the government or a private institution, or a reimbursement from an employer for training expenses, should be subtracted from the deducted amount. Over the period 2006 – 2012, a threshold of 500 euro applied to all deductible lifelong learning expenditures in a given year. The maximum deductible amount each year was (and is) 15,000 euro.

The deductible for lifelong learning expenditures changed quite substantially in 2013. First, the threshold was reduced from 500 euro to 250 euro. Second, the deductible became limited to tuition fees and compulsory additional learning tools, such as books and protection materials. This meant for example that the depreciation of a computer was no longer deductible. These changes provide another reason why we limit ourselves to the 2006 – 2012 period in the main analyses.

While training expenditures are typically individual expenditures, partners can choose whether they deduct the expenditures from their own taxable income or whether they transfer the expenditures to their partner who can then subtract it from his or her taxable income. To minimize the household tax burden, partners typically shift the tax deduction to the partner that has the higher marginal tax rate (see Section 5.5). The threshold of 500 euro must first be applied to each partner’s personal expenditures before the expenditures can be shifted between partners. For couples we use data on personal or 'own' expenditures and data on declared expenditures to show the importance of accounting for this shifting behaviour.

5.3 Theoretical framework

Following Leuven and Oosterbeek (2012), we illustrate the basic mechanism via which a tax deduction for lifelong learning expenditures in combination with differences in marginal tax rates affects the investment in lifelong learning in a stylized life-cycle model.

Lifetime utility depends on consumption in period 1 and 2: $U(C_1, C_2)$. We assume that the utility function is additively separable in period 1 utility and period 2 utility, and period 2 utility is discounted by a factor $1/(1+\delta)$, where δ is the subjective discount rate:

$$U(C_1, C_2) = U(C_1) + \frac{1}{1+\delta} U(C_2). \tag{5.1}$$

Consumption in period 1 depends on gross income w_1, lifelong learning expenditures L, the tax rate τ_1 and savings S:

$$C_1 = (1 - \tau_1)(w_1 - L) - S, \tag{5.2}$$

noting that lifelong learning expenditures are deducted from gross income rather than net income. Also note that for simplicity we assume that agents face a flat tax system. Consumption in period 2 then depends on gross income w_2, the return on lifelong learning expenditures, the tax rate τ_2 and the return on period 1 savings:

$$C_2 = (1 - \tau_2)(w_2 + f(L)) + (1 + r)S, \tag{5.3}$$

where $f(L)$ is the return on lifelong learning expenditures in terms of a higher gross period 2 income, for which we assume $f(0) = 0$, $f' > 0$ and $f'' < 0$, and r is the return on savings.

Maximizing lifetime utility with respect to lifelong learning expenditures and savings gives, respectively:

$$\frac{\partial U(.)}{\partial L} = 0 \Rightarrow U'_{C_1}(-(1-\tau_1)) + \frac{1}{1+\delta} U'_{C_2}(1-\tau_2) f'(L) = 0, \tag{5.4}$$

$$\frac{\partial U(.)}{\partial S} = 0 \Rightarrow U'_{C_1}(-1) + \frac{1}{1+\delta} U'_{C_2}(1+r) = 0. \tag{5.5}$$

Solving for L then gives the implicit function:

$$f'(L) = \frac{(1-\tau_1)}{(1-\tau_2)}\frac{(1+r)}{(1+\delta)}. \tag{5.6}$$

In the empirical application below we will compare individuals with a lower τ_1, with a taxable income just below a tax bracket threshold, with individuals with a higher τ_1, with a taxable income just above a tax bracket threshold. Equation (5.6) shows that *ceteris paribus*, individuals with a higher τ_1 will invest more in lifelong learning than individuals with a lower τ_1. Indeed, when τ_1 is higher, the right hand side of (5.6) is lower. Hence, at the optimum, $f'(L)$ will be lower as well, and given that $f''(L) < 0$, this implies that L should be higher. Intuitively, the investment cost of lifelong learning is lower when τ_1 is higher. In the appendix we show that ceteris is indeed very close to paribus as individuals just below and just above income tax bracket thresholds are very similar in observable characteristics (and hence in r and δ in terms of our simple stylized model).[7] They also face very similar tax rates τ_2 in years after the lifelong learning investment, see Figure A4.[8]

5.4 Empirical methodology

We apply a different empirical methodology for singles and couples. The tax deduction introduces a kink in the effective costs of lifelong learning expenditures. Therefore we prefer to use a regression kink design, provided that the conditions for using a regression kink design hold.[9] A crucial condition for a regression kink design is that there is no bunching around the kink. Below we show that this condition holds for singles, but not for couples. As discussed in Section 5.2, couples can shift their lifelong learning expenditures between partners. Couples who minimize their joint tax burden will generally shift deductibles to the partner with the highest marginal tax rate, which will typically be the highest earning partner, until marginal tax rates are equal. This means that the highest earner often ends up close to the beginning of a tax bracket. This creates bunching at the kink, which invalidates the assumptions

[7] See Figure A5 for singles. This does not hold for couples due to shifting of deductibles (see Figure A6). We discuss our strategy to deal with this in section 5.4.2. Note that in terms of gross income couples just below and above the tax brackets are very similar (Figure A10).

[8] Note that when $\tau_1 = \tau_2$, lifelong learning expenditures do not depend on marginal tax rates (Boskin, 1975; Eaton and Rosen, 1980; Leuven and Oosterbeek, 2012). However, below we show that this does not hold for large parts of the individuals in the sample. Indeed, the analysis rests on the fact that τ_1 is different just below and above tax bracket thresholds, whereas τ_2 is very similar.

[9] See Card et al. (2015a), Card et al. (2015b) and Landais (2015) for an introduction to the regression kink design methodology.

underlying the regression kink design. For couples we therefore use a so-called donut regression discontinuity design.[10] In the donut regression discontinuity design we drop observations from income bins around the kink for which we observe excess mass. The size of the donut in our preferred specification (1,000 euro on either side of the kink) is so large that for the large majority of the sample to the right of the kink included in the regression there is a fixed difference, or discontinuity (as opposed to a kink), in the financial gain from the tax deduction. Therefore, we apply a donut regression discontinuity design for couples.

5.4.1 Singles: regression kink design

For singles we exploit the differences in the marginal tax rates in a regression kink design to identify the causal effect of the tax deduction on lifelong learning expenditures. The general idea is that the outcome variable is a continuous function of income in the absence of the tax deduction, but that the tax deduction in combination with a discontinuity in the marginal tax rate creates an exogenous kink in the effective costs of lifelong learning, which potentially results in a kink in the use of and expenditures on lifelong learning as well.

Figure 5.1 illustrates the kink when going from the third to the fourth bracket, located at a taxable income of 52,000 euro. Suppose that an individual has 2,500 euro lifelong learning expenditures. The marginal tax rate to the left of the kink is 42%. The effective costs of the lifelong learning expenditures to the left of the kink then are $(1 - 0.42) * (2,500 - 500) + 500 = 1,660$ euro. When the individual has taxable income (before the tax deduction is applied) in the fourth tax bracket, the effective costs of lifelong learning expenditures are lower. For example, at 1,000 euro to the right of the threshold, the effective costs of lifelong learning are $(1-0.52)*(2,500-1,500)+(1-0.42)*(1,500-500)+500 = 1,560$ euro, or 6% less than on the left-hand side of the threshold. Finally, for individuals with a taxable income 2,000 euro to the right of the threshold and beyond, the effective costs of lifelong learning are $(1 - 0.52) * (2,500 - 500) + 500 = 1,460$ euro, or 12% less than on the left-hand side of the threshold. This suggests running a regression kink design using observations up to the point where the financial gain becomes constant again, and for symmetry we then also use observations from the same distance to the kink on the left hand side. Hence, using observations from the interval [50,000,54,000], indicated by the dashed lines in Figure 5.1.

[10]See Barreca et al. (2011), Barreca et al. (2016) and Hoxby and Bulman (2016) for an introduction to the donut regression discontinuity design methodology.

Figure 5.1: Effective costs of 2,500 euro lifelong learning expenditures

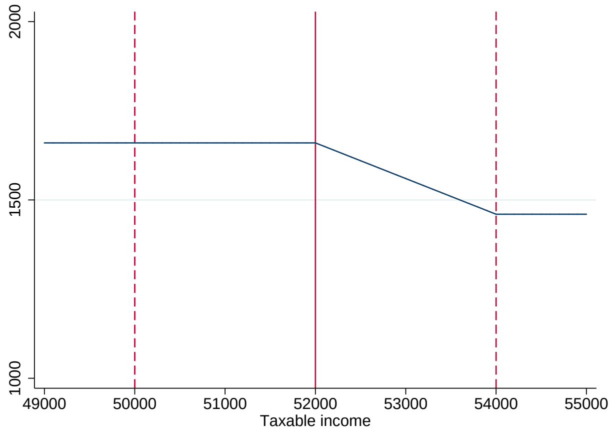

We estimate the effect of the tax deduction on i) the probablity of filing lifelong learning expenditures, and ii) the amount of lifelong learning expenditures filed (including the zeros), using the following linear model:[11]

$$Y_{it} = \alpha + \beta R_{it} + \delta 1(R_{it} > 0) * R_{it} + \gamma X_{it} + \eta_t + \epsilon_{it}, \tag{5.7}$$

where i denotes the individual and t denotes calender year. R_{it} is (recentered) taxable income before deducting lifelong learning expenditures, the parameter δ measures the treatment effect, the change in the slope at the kink. X_{it} are a set of demographic control variables, η_t are year fixed effects and ϵ_{it} is the error term. To account for correlation in the error term at a level higher than the individual, we use cluster-robust standard errors for income groups of 100 euro (Bertrand et al., 2004; Donald and Lang, 2007).[12]

5.4.2 Couples: donut regression discontinuity design

Couples can manipulate their taxable income by shifting deductibles between fiscal partners, including but not limited to the deduction for lifelong learning expenses. In the empirical analyses we show that we indeed observe bunching at the cutoff for couples.[13] To mitigate this problem we apply a so-called donut regression discontinuity design, where we drop observations around the cutoff. We present

[11] For the probability of filing lifelong learning expenditures this is a linear probability model (Angrist and Pischke, 2009).

[12] Standard errors are very similar when we do not use cluster-robust standard errors, see Table A4 in the appendix.

[13] Recall that we measure income before the deduction for lifelong learning expenditures is subtracted. Hence, lifelong learning expenditures do not cause the bunching we observe in the data.

results for various sizes of the donut hole, including no donut hole as in the standard regression discontinuity setup.

As discussed above, by applying a large donut hole, we are basically left with a discontinuity in the effective costs of schooling between those on the left and right hand side of the donut hole. This means that for couples the 'treatment effect' is measured for a discontinuity, where we compare those to the right of the donut hole with those to the left. We therefore estimate the following regression discontinuity model excluding the observations close to the threshold:

$$Y_{it} = \alpha + \beta R_{it} + \gamma 1(R_{it} > 0)R_{it} + \delta 1(R_{it} > 0) + \phi X_{it} + \eta_t + \epsilon_{it}, \qquad (5.8)$$

where most terms are defined as above. The treatment effect δ however, is now measured by the change in the intercept to the right of the threshold. Also for the donut regression discontinuity design we use cluster-robust standard errors, clustered at income groups of 100 euro.

5.5 Data

For the empirical analysis we use the universe of Dutch tax payers, available via the remote access server of Statistics Netherlands. We have data for the period 2006–2013, but we focus on the period 2006–2012. During the period 2006 – 2012 the tax deduction for lifelong learning expenditures remained largely unchanged.

We make the following selections. We drop all individuals younger than 25 years of age or older than 60 years of age. Furthermore, we drop individuals who are enrolled at a full-time higher education institution, because students can use the tax deduction for other reasons than lifelong learning expenditures. We also exclude individuals on retirement benefits, on other types of social insurance and individuals without income, because their demographic characteristics are quite different from the rest of the sample. Finally, for couples we only keep those where both partners are still in the sample after we made the selections above.

As dependent variables we consider the take-up rate of the lifelong learning tax deduction and the deducted amount. We subtract the threshold of 500 euro from the deducted amount before we calculate the take-up rate (dummy) and the deducted amount.

The bunching is caused by other deductibles that can be shifted between partners and that have already been deducted from the income definition that we use. Specifically, our running variable is taxable individual income plus the deduction for lifelong learning expenditures. Individual gross incomes show no bunching around the kinks (see Figure A9 in the appendix).

Couples can shift the deductible amount from one partner to the other. When the marginal tax rates differ, the household will be better off financially when the partner with the lower marginal tax rate shifts the lifelong learning expenditures to the partner with the higher marginal tax rate. Indeed, this is what most couples do, see Table 5.2. Close to 83% of people with a lower marginal tax rate than their partner shift the lifelong learning expenditures to the partner with a higher marginal tax rate. Therefore, for couples it is important to distinguish between what we denote as the own deducted amount and the declared amount, where the latter includes the amount (above the threshold) coming from or going to the other partner (hence the declared amount can be higher or lower than the own amount).[14]

Table 5.3 shows the distribution of the use of the deductible by income level. For singles, 31% of the population has taxable income below 20,000 euros, and about 2.6% uses the deductible in 2006–2012. For singles with an income between 20,000 and 40,000 euros, the largest group, about 3.1% use the deductible. Higher income singles make up a much smaller share of the population, but are more likely to use the deductible. In addition to their more frequent use, they also deduct higher amounts. Especially singles with a taxable income of more than 60,000 euros – about 4% of the population of singles – have a high deductible at close to 3,500 euros for those using the deductible. For couples we find similar patterns, with higher earners both more likely to use the deductible and deducting higher amounts.

We study two discontinuities in marginal tax rates: 1) the increase in the marginal tax rate when we move from the first to the second tax bracket, which we indicate with 'kink 1', and 2) the increase in the marginal tax rate when we move from the third to the fourth tax bracket, which we indicate with 'kink 2'.

Descriptive statistics for singles are given in Table 5.4. In the first column we present descriptive statistics for the sample around kink 1. Specifically, these are statistics for the sample in our preferred specification with individuals from −1,330 to +1,330 euro around kink 1. 2.9% of this sample deducts lifelong learning expenditures, and the average amount deducted is almost 40 euro (including the zeros). The average amount is 1,330 euro per person that uses the deduction, which motivates the sample interval that we use. 66% of the sample around kink 1 are female, they are on average 40 years of age, have 0.8 children on average and 15% of them are born outside the Netherlands or has at least one parent born outside the Netherlands ('Foreign'). We have about 660,000 observations in this sample.

[14]Typically the declared amount will be higher than the own amount for primary earners and lower than the own amount for secondary earners.

Table 5.2: Shifting of lifelong learning expenditures in couples (in %)

Marginal tax rate relative to partner	No shifting	Partial shifting	Full shifting	Total
Higher	89.7	8.8	1.6	100
Equivalent	54.0	22.2	23.7	100
Lower	7.6	9.6	82.8	100

Notes: Own calculations based on tax return data from Statistics Netherlands.

Table 5.3: The distribution of the use of the deductible

Taxable income	Share in population	Share using the deductible	Avg deductible	Avg deductible for users
A. Singles				
$< 20,000$	0.31	0.026	33.87	$1,287.25$
$20,000 - 39,999$	0.53	0.031	43.45	$1,389.30$
$40,000 - 59,999$	0.12	0.036	67.58	$1,882.47$
$> 60,000$	0.04	0.033	113.20	$3,471.08$
B. Couples				
$< 20,000$	0.37	0.008	7.13	866.63
$20,000 - 39,999$	0.42	0.025	29.18	$1,170.31$
$40,000 - 59,999$	0.14	0.030	45.05	$1,486.62$
$> 60,000$	0.07	0.034	64.83	$1,926.51$

Notes: Full population between 20 and 60 years old for 2006–2012. Number of observations: 30,191,548. Incomes and amounts in real 2012 euros.

The second column gives the descriptive statistics for the sample around kink 2 for our preferred specification with a bandwidth of 2,000 euros around the kink. The take-up rate is higher for this group, 3.9%, and the average amount is also higher at around 81 euro (including the zeros, the average amount is 2,091 euro per person that uses the deduction).

There are fewer females in the sample around kink 2 (32%), on average they are somewhat older, have fewer children and are less likely to be from foreign parents than around kink 1. This sample is smaller, with close to 200,000 observations. These individuals are already relatively high in the income distribution (approximately 10% of the population with income has income in the fourth (top) bracket in the Netherlands).

Descriptive statistics for couples are given in Table 5.5. We now present statistics for the preferred sample using a bandwidth of 5,000 euro to the left and right of the kink and applying a donut hole of 1,000 euros to the left and to the right of the kink. We present descriptives separately for primary and secondary earners. 2.6% of primary earners around kink 1 declares lifelong learning expenditures, and on average they declare 31 euro (1,208 euro per declaring person). The percentage of primary earners declaring own lifelong learning expenditures is substantially lower at 1.8%, and also the average amount is substantially lower at 21.9 euro (1,217 euro per declaring person). Turning to the demographic control variables, only 28% of these primary earners around kink 1 are female, the average age is close to 39 years of age, they have 1.3 children on average and only one in ten is foreign born or has one or more foreign born parents.

Secondary earners around kink 1 are less likely to declare lifelong learning expenditures, only 1.4%, and on average they declare 13 euro (929 euro per declaring person). However, the percentage of secondary earners declaring own lifelong learning expenditures is actually somewhat higher than for primary earners, 1.9%, and also the average amount is somewhat higher at 22.4 euro (1,171 euro per declaring person). Secondary earners are more likely to be female, they are on average about a year younger than the primary earners, have the same number of children and are about equally likely to be foreign born or have one or more foreign born parents. We have about half a million couples in the sample of our preferred specification for couples around kink 1.

Moving to kink 2, we observe a much higher share of primary earners declaring lifelong learning expenditures, 3.8%, at an average amount of 62 euro (1,635 euro per declaring person). However, they are much less likely to declare own lifelong learning expenditures, 2.1%, and also the average own amount of 36.7 euro is much smaller (1,722 euro per declaring person). The large majority of these primary earners are male, they are older than at kink 1, have about the same number of children and are much less likely to be foreign born or have a foreign born parent. For secondary earners we again see a much lower share declaring lifelong learning expenditures, 1.3%, with an average amount of 16 euro (1,244 euro per declaring person). However, the share of secondary earners declaring own lifelong learning expenditures is again higher than for primary earners, 2.8%, with an average amount of 42 euro (1,466 euro per declaring person). These secondary earners are predominantly female, are one and a half year younger than the primary earners on average, have the same number of children, and are also not very likely to be foreign born or have one or more foreign born parents. For couples around kink 2 we have about three quarters of a million observations.[15]

[15] Exploiting the panel nature of our data we also examined persistence in the use of the deductible. We find that close to 50% of the people who use the deductible in year t have also used it in year $t-1$, and around 30% also used it in year $t-2$. After controlling for observables, using the deductible in year $t-1$ ($t-2$) increases the probability of using the deductible in year t by 38% (13%). Hence, there seems to be substantial persistence in the use of the deductible. This

Table 5.4: Descriptive statistics: singles

	Kink 1	Kink 2
Outcome variables		
Deductible	0.0292	0.0390
	(0.1684)	(0.1936)
Deducted amount	39.2106	81.3248
	(346.5883)	(800.5169)
Control variables		
Female	0.6565	0.3194
	(0.4749)	(0.4662)
Age	39.9325	43.5710
	(9.8179)	(9.1946)
Number of children	0.8397	0.4886
	(0.9849)	(0.8523)
Foreign	0.1470	0.0583
	(0.3542)	(0.2343)
Observations	663,368	197,584

Notes: Sample period 2006–2012. Standard deviations reported in parentheses. Descriptives are presented for the preferred sample using a bandwidth of 1,330 euros for kink 1 and 2,000 euros for kink 2. The deducted amount includes the zeros of non-users.

Table 5.5: Descriptive statistics: couples

	Kink 1		Kink 2	
	Primary earner	Secondary earner	Primary earner	Secondary earner
Outcome variables				
Declared deductible	0.0260	0.0138	0.0381	0.0130
	(0.1591)	(0.1168)	(0.1914)	(0.1131)
Declared deducted amount	31.3919	12.8519	62.2872	16.1188
	(276.3430)	(168.7083)	(503.0561)	(229.1265)
Own deductible	0.0180	0.0191	0.0213	0.0284
	(21.8619)	(0.1368)	(0.1444)	(0.1662)
Own deducted amount	21.8619	22.3547	36.7184	41.6877
	(168.7083)	(233.5227)	(417.4154)	(367.8090)
Control variables				
Female	0.2787	0.7255	0.1118	0.8868
	(0.4483)	(0.4463)	(0.3152)	(0.3169)
Age	38.7043	37.6005	44.9257	43.2654
	(8.5814)	(8.4261)	(8.0111)	(8.0296)
Number of children	1.3111	1.3111	1.4321	1.4321
	(1.0226)	(1.0226)	(1.0761)	(1.0761)
Foreign	0.1184	0.1207	0.0265	0.036
	(0.3231)	(0.3258)	(0.1605)	(0.1864)
Observations	498,627	498,627	756,617	756,617

Notes: Sample period 2006–2012. Standard deviations reported in parentheses. Descriptives are presented for the baseline sample with a 1000 euro donut hole. The deducted amounts include the zeros for non-users.

could be because some courses simply take multiple years to complete. Another reason could be that once people are aware of the deductible, they tend to use it more frequently.

5.6 Results

5.6.1 Singles

First we consider the results for singles. Figure 5.2a and 5.2b present graphical evidence for bunching (or heaping) around kink 1 and 2 respectively, and hence the potential role of manipulation of the running variable. On the horizontal axis we have taxable income plus the declared lifelong learning expenditures (potentially zero) relative to the kink, using bins of 100 euro. On the vertical axis we have the density. At both kinks there is no clear evidence of bunching. If anything there appears to be some excess mass only at the first bins of 100 euro next to the kink.[16] This suggests that singles essentially do not manipulate their income relative to these kinks.[17]

Figure 5.2c and 5.2d show the take-up rate of the deductible for schooling expenditures (in excess of the minimum expenditure threshold) for kink 1 and 2, respectively. We present averages per bin by income. The solid red lines give the predicted take-up rate, using a linear regression without demographic control variables, allowing for a different slope to the right of the kink. The dashed red lines give the corresponding 95% confidence intervals.

Above the graph we report the corresponding coefficient for the change in the slope on the right-hand side. The graph and the estimated coefficient suggest zero effect for kink 1, but a positive and statistically significant effect for kink 2. Figure 5.2e and 5.2f plot the declared amount of schooling expenditures for singles around kink 1 and kink 2 (above the threshold, and including the zeros). Again, there is no apparent kink in the relation between the declared amount and taxable income at kink 1, but there appears to be a kink, albeit not statistically significantly different from zero, in the relation between the declared amount and taxable income at kink 2. Furthermore, for kink 2, we also see a 'flattening out' of the effect on the take-up

[16] Following McCrary (2008) and Cattaneo et al. (2017), we study the excess mass using a density test. For kink 2 this gives a p-value for the null hypotheses of no excess mass of 0.21, 0.41 and 0.71 using the conventional, undersmoothed and robust-bias corrected of the Stata package *rddensity*, respectively. The conventional approach may be asymptotically biased. The undersmoothed and robust-bias corrected approaches try to correct for this bias in different ways, see Cattaneo et al. (2016, 2017) for more detail. For kink 1 the p-values for the different methods are 0.07, 0.07 and 0.02 (the latter suggests that there might be some excess mass at kink 1, but Figure 5.2a shows that the excess mass is small and very local). Furthermore, there are no discontinuities in the demographic control variables around kink 1 or 2 for singles, see Figure A5.

[17] Empirical studies looking at bunching at tax bracket thresholds typically find little evidence of bunching, at least for wage earners, see e.g. Kleven (2016) for an overview.

Figure 5.2: Density around the kink, probability to use the deductible and the deducted amount: singles

Kink 1 | Kink 2

Density around the cutoff

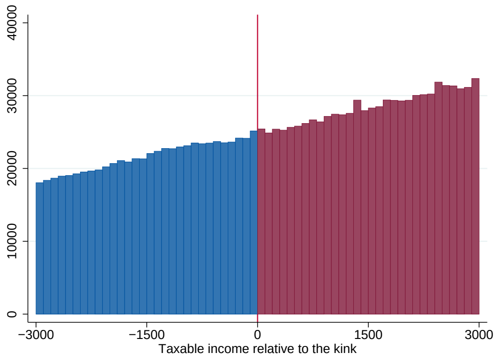

(a)

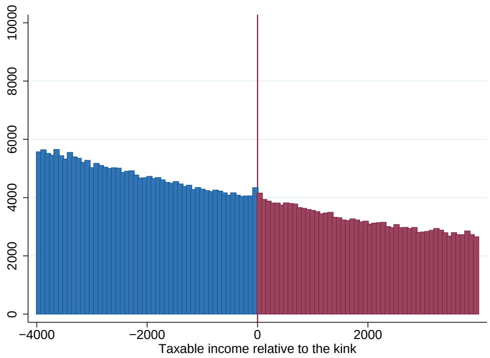

(b)

Take-up rate of the deduction

Estimate of the discontinuity: −0.0013 (0.0011)

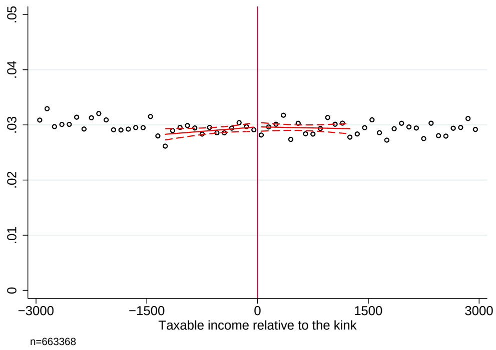

(c)

Estimate of the discontinuity: 0.0037 (0.0015)**

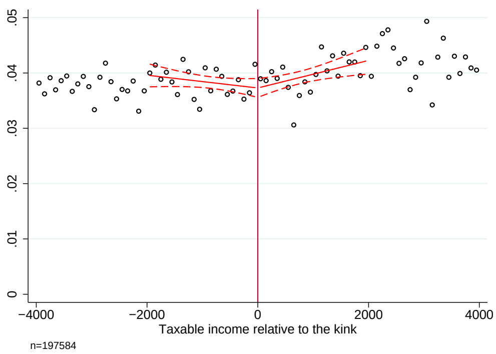

(d)

Deducted amount

Estimate of the discontinuity: −1.9414 (2.2131)

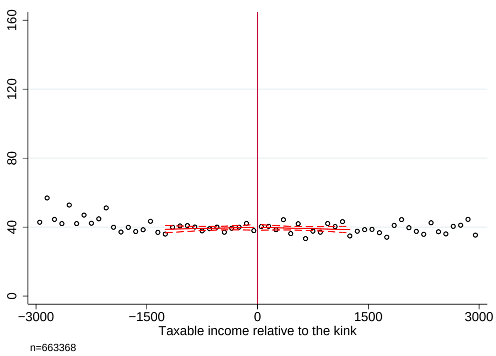

(e)

Estimate of the discontinuity: 5.2149 (6.2460)

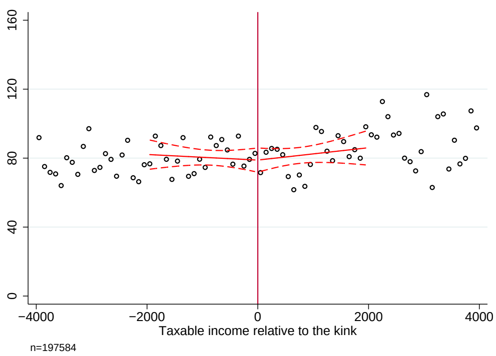

(f)

Notes: Own calculations based on register data from Statistics Netherlands. The regression lines are linear functions without any control variables, with a separate intercept and slope on the right-hand side of the kink. The deducted amount includes the zeros for non-users. Estimates for kink 1 include observations from minus 1,330 to plus 1,330 euro relative to the kink. Estimates for kink 2 include observations from minus 2,000 to plus 2,000 euro relative to the kink.

rate and the deducted amount, which is consistent with the flattening out of the financial gain to the right of the kink (see Section 5.4).

However, this is not controlling for demographic characteristics. Our simple theoretical model suggests that it could be important to control for observable characteristics. In Panel A in Table 5.6, we present regression results for the regression-kink coefficient (change in the slope) without and with demographic control variables and for different bandwidths. Column (1) gives the results for the probability of using the lifelong learning deduction without demographic control variables. For all bandwidths we find a small and statistically insignificant effect. The results are very similar when we include demographic control variables in Column (2). Our preferred specification includes demographic control variables and uses a bandwidth of 1,330 euro. Here we find an effect of -0.0014. The running variable is in thousands of euro, hence the interpretation is that the additional financial gain of having an income 1,000 euro to the right of the kink, leads to a (counterintuitive) drop in the take-up rate of the lifelong learning deduction of -0.14 percentage points, but as noted above the effect is not statistically significantly different from zero. Our preferred bandwidth is 1,330 euro because this is the average amount of schooling expenditures deducted at 1,330 euro to the right of the kink, which is where the kink ends on average.[18] Also for the deducted amount we find a small and insignificant (negative) effect, with and without demographic control variables, see Column (3) and (4), respectively.

Panel B in Table 5.6 gives the regression results for the regression-kink coefficient for kink 2, again without and with demographic control variables and for different bandwidths. For kink 2 our preferred bandwidth is 2,000 euro, which is very close to the average lifelong learning expenditures deducted to the right of the kink of 2,060 euro at 2,060 euro. For this bandwidth we find a statistically significant positive effect of 0.38 percentage points. A bandwidth that is somewhat smaller or larger results in a somewhat lower coefficient, but not statistically significantly different from our preferred bandwidth.

The regression results for the average deducted amount for different bandwidths for kink 2 are given in Column (3) and (4) of panel B, without and with demopgrahic control variables, respectively. Again, accounting for demographic control variables

[18]Figure A8 in the appendix shows that the average amount of schooling expenditures is rather stable over income bins. Furthermore, we do not exploit the ‘second kink’ to the right of kink 1, where the financial gain is no longer increasing on average, because the exact location of this ‘second kink’ depends on the individual amount of lifelong learning expenditures, which varies across individuals with the same income.

Table 5.6: Treatment effect estimates for singles on the probability to use the deductible and the deducted amount, for different bandwidths around the kink

	(1)	(2)	(3)	(4)	
	Use of the deductible		Deducted amount (in euro)		Observations
	No controls	Controls	No controls	Controls	
Panel A: Kink 1					
Bandwidth					
1,000	0.0006	0.0007	−2.2945	−2.3580	496,957
	(0.0015)	(0.0016)	(3.5096)	(3.5187)	
1,330	−0.0014	−0.0014	−2.4042	−2.4376	662,848
	(0.0012)	(0.0012)	(2.4114)	(2.3543)	
1,500	−0.0006	−0.0006	−1.4232	−1.5225	749,526
	(0.0011)	(0.0011)	(2.0498)	(2.0209)	
Panel B: Kink 2					
Bandwidth					
1,500	0.0024	0.0024	4.5190	4.8451	148,526
	(0.0021)	(0.0020)	(6.5772)	(6.5462)	
2,000	0.0038***	0.0038***	5.5721	5.8728	197,584
	(0.0011)	(0.0010)	(3.9432)	(3.8610)	
2,500	0.0031***	0.0032***	7.8314**	8.0554**	246,949
	(0.0009)	(0.0009)	(3.3579)	(3.3800)	

Notes: Sample period 2006–2012. Cluster-robust standard errors clustered by income bins of 100 euro in parentheses, *** p<0.01, ** p<0.05, * p<0.1. All regressions include year fixed effects. The regressions with controls include gender, ethnicity, age, age^2 and the number of children in the household as demographic controls. Full estimation results for our preferred specification with a bandwidth of 1,330 euro for Kink 1 and 2,000 euro for Kink 2 are reported in Table A3 in the appendix.

hardly affects the results. For our preferred bandwidth of 2,000 euro, and including demographic control variables, we find a positive coefficient of 5.9 euro. However, this coefficient is not statistically significant.

A number of robustness checks are given in the appendix. Results without clustering standard errors are reported in Table A4 in the appendix. This results in slightly larger standard errors. The estimated treatment effect for our preferred bandwidth of 2,000 euro around kink 2 remains significant at the 5% level, compared to 1% in the baseline.

To further investigate the robustness of our estimates we also performed the permutation test outlined in Ganong and Jäger (2018). We construct a distribution of placebo estimates in regions where we know there is no kink in the tax system and regions where we know there is a kink. Using our preferred bandwidth of 1,330 euros, we find no statistically significant treatment effect at kink 1.[19] For kink 2, using

[19]We imposed placebo kinks using data ranging from $[-2,000; 2000]$ around the actual kink up to $[-8,000; 8,000]$, excluding data around the actual kink. The p-values for no treatment effect at kink 1 range from 0.68 to 0.96.

our preferred bandwidth of 2,000 euros, we do find a highly statistically significant treatment effect.[20] This confirms our main results.

Finally, for our baseline estimates we use a linear specification, separately estimated for data on the left and on the right of the kink. We have also estimated quadratic and cubic specifications. For kink 1 we find no evidence for a kink using these more flexible specifications. For kink 2 and using our preferred bandwidth, we find smaller estimates that are no longer statistically significant. However, if we use a slightly larger bandwidth of 2,500 euros the treatment again becomes statistically significant for both the quadratic and cubic specification, suggesting that a lack of statistical power may play a role here.

Table 5.7 gives the estimated treatment effect for subgroups of singles, for subgroups that differ in their demographic characteristics.[21] We focus on the effect on the take-up rate of the deductible around kink 2. The treatment effects around kink 1 and for the deducted amount (both at kink 1 and 2) are typically small and insignificant.[22] We find that the treatment effect is somewhat larger for men than for women, substantially larger for singles whose parents were born in the Netherlands than for singles with one or more foreign parents and also substantially larger for high-educated than for low-educated singles (note that this is conditional on having income around kink 2). Furthermore, we find that the treatment effect is relatively large for middle-aged persons, 35–39 and in particular 40–44 years of age, compared to younger and older workers.[23] Indeed, middle-age may be the time to invest in another job, whereas skills are probably more up-to-date for young workers and the return period for investments in work-related human capital is rather short for older workers.

Returning to our base estimate for all singles, we can convert our preferred estimate to an elasticity of the probability of (deducting) lifelong learning expenditures with respect to the effective costs of lifelong learning expenditures. Consider an individual that has 2,500 euro in lifelong learning expenditures, or 2,000 euro above the threshold (which is close to the average around kink 2). Furthermore, suppose that this individual has an income that is 1,000 euro to the right of the kink, which is in the middle of the region where the financial gain increases. For this

[20] We imposed placebo kinks using data ranging from $[-5,000; 5000]$ to $[-10,000; 10,000]$ around the actual kink, excluding up to 4,000 euros left and right of the actual kink. We find p-values for the kink ranging from 0.0325 to 0.

[21] For all subgroups we use the same specification and bandwidth as in the baseline for all singles.

[22] Available on request.

[23] Note that the treatment effects for subgroups are not statistically significantly different from each other.

Table 5.7: Treatment effect by demographic characteristics for singles at kink 2

	Use of the deductible	Observations
All (base)	0.0038***	197,584
	(0.0010)	
Women	0.0032	63,100
	(0.0027)	
Men	0.0040***	134,484
	(0.0013)	
Native	0.0043***	164,254
	(0.0012)	
Foreign	0.0014	33,330
	(0.0028)	
Low educated	0.0029	3,872
	(0.0052)	
Middle educated	0.0061	26,714
	(0.0045)	
High educated	0.0078***	72,292
	(0.0028)	
25 - 29 years	−0.0011	17,610
	(0.0069)	
30 - 34 years	0.0069	32,317
	(0.0045)	
35 - 39 years	0.0067*	36,869
	(0.0038)	
40 - 44 years	0.0083***	38,668
	(0.0025)	
45 - 49 years	0.0019	39,004
	(0.0020)	
50 - 54 years	−0.0020	36,118
	(0.0033)	
55 - 60 years	0.0033	29,983
	(0.0021)	

Notes: Sample period 2006–2012. Cluster-robust standard errors clustered by income bins of 100 euro in parentheses, *** p<0.01, ** p<0.05, * p<0.1. All regressions are separately estimated for each subgroup and include year fixed effects and controls for the other demographic characteristics.

individual we predict an increase in the take-up rate of 0.38 percentage points, or about +10% relative to the baseline of 3.8 percentage points left of the kink. The effective costs of lifelong learning are about 6% lower than to the left of the kink, see Section 4. The elasticity of the take-up rate of (deducting) lifelong learning expenditures with respect to the effective costs of lifelong learning expenditures is then $+10\%/(-6\%) \approx -1.7$ with a 95% confidence interval of $[-0.8, -2.5]$.[24]

[24]In a similar way, we can calculate the elasticity of the amount of lifelong learning expenditures with respect to the effective costs of lifeling learning expenditures. For an individual that has an income 1,000 euro to the right of the kink we predict an increase in (deducted) lifelong learning expenditures of 5.9 euro, or about +7% relative to the baseline of 79 euro to the left of the kink. Relating this to the drop in the effective costs of lifelong learning expenditures of −6%, the elasticity of (deducted) lifelong learning expenditures with respect to the effective costs of lifelong learning expenditures in then $+7\%/(-6\%) \approx -1.2$ with a 95% confidence interval of $[-2.8, 0.4]$

5.6.2 Couples

Next, we consider the effects for couples. Within couples, we study the effects for the partners with the highest gross income in the household or 'primary earners' and the effects for the partners with the lowest gross income in the household or 'secondary earners'. Furthermore, we present results for the declared amount and the own amount. Because partners can shift the lifelong learning expenditures from one partner to the other when they file their taxes, the effect on the declared amount and own amount can differ. Indeed, we show that this makes a big difference to the results, which underscores the value of the rich data that we use in the analysis.

Figure 5.3a and 5.3b present graphical evidence on the role of bunching of taxable income plus the lifelong learning deduction for primary earners around kink 1 and kink 2, respectively. The figures provide clear evidence of bunching, and hence of manipulation of the running variable near the kinks.[25] Indeed, by shifting deductibles (other than the lifelong learning expenditures deduction) from the secondary earner to the primary earner, couples can reduce the tax burden of the household, up to the point where the marginal tax rate of the primary earner is no longer higher than the marginal tax rate of the secondary earner.[26] This bunching will generate a bias in the estimate when couples that are more likely to use the lifelong learning tax deduction are also more likely to manipulate their income. This is why we apply a donut hole to the sample included in the estimates in our preferred specification. Furthermore, to have enough observations we include households with a running variable plus and minus 5 thousand euro of the kink. For the large majority of couples included in the sample to the right of the kink there is then a fixed difference between the financial gain on the left-hand and on the right-hand side, i.e. a discontinuity rather than a kink. Therefore, we estimate the effect using a donut regression discontinuity design.

For couples we prefer to look at their *own* (declared) lifelong learning expenditures rather than their total declared expenditures. When using the total declared schooling expenditures, the treatment effect consists of the effect on the own lifelong learning expenditures by primary earners and the shifting of lifelong learning expenditures

[25]The p-values for the McCrary density tests of no excess mass are all below 0.0001.

[26]The RD plots of the demographic control variables for primary earners also show discontinuous jumps around kink 1 and kink 2, see Figure A6 in the appendix, again suggesting manipulation of the running variable. Figure A9 in the appendix shows that there is no bunching around the kink if we use gross income instead of taxable income, the income before applying any of the deductibles. Figure A10 confirms that there are also no discontinuities in observable characteristics of primary earners at the kink if we use gross income instead of taxable income. This shows that the bunching that we find in taxable income (plus the lifelong learning tax deduction) is due to shifting of deductibles (other than the lifelong learning tax deduction).

Figure 5.3: Own use of the deductible and own amount: primary earners

Kink 1 | Kink 2

Density around the cutoff

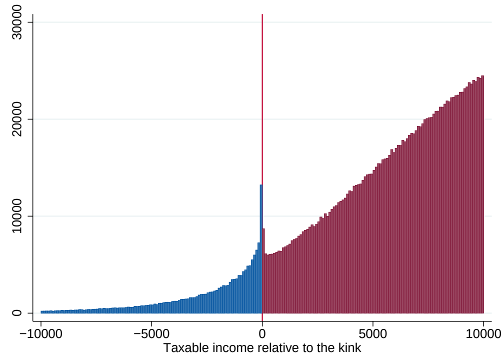

(a)

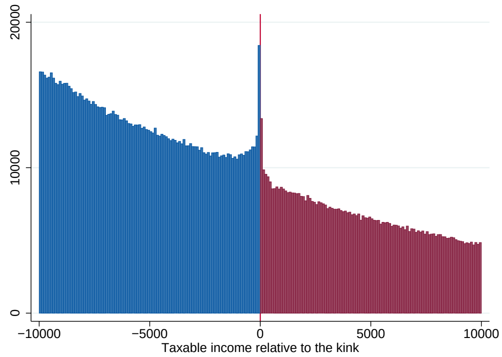

(b)

Take-up rate of own deductible

Estimate of the discontinuity: 0.0042
(0.0013)***

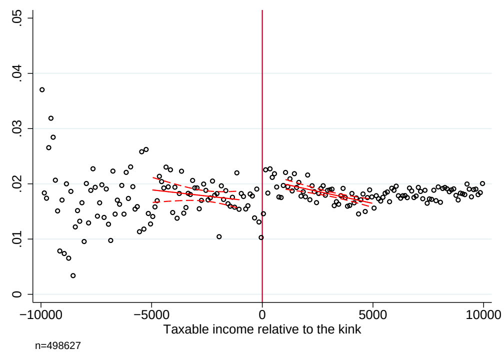

(c)

Estimate of the discontinuity: 0.0018
(0.0009)**

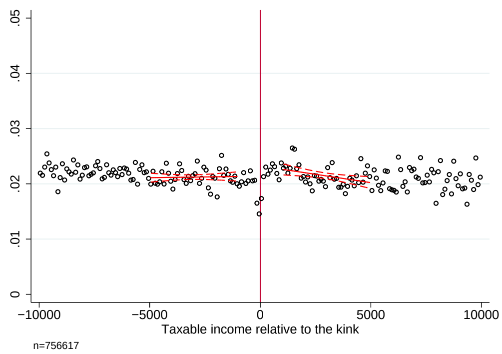

(d)

Own deducted amount

Estimate of the discontinuity: 6.3545
(2.3944)***

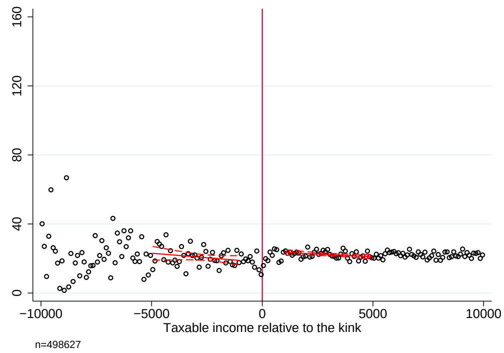

(e)

Estimate of the discontinuity: 7.8845
(2.7079)***

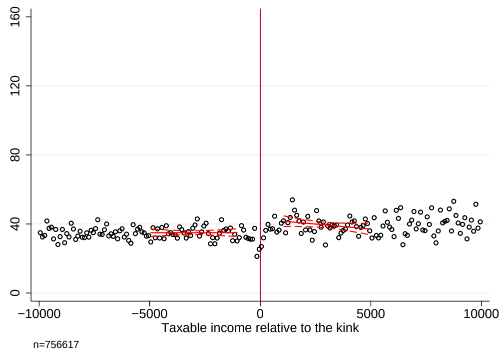

(f)

Notes: Own calculations based on register data from Statistics Netherlands. The regression lines are linear functions without any control variables, with a separate intercept and slope on the right-hand side of the kink taking into account a donut hole of 1,000 euro. The estimate of the discontinuity at the cutoff presented above each figure takes into account the donut hole.

from the secondary earner to the primary earner. The second effect does not reflect an actual increase in lifelong learning expenditures, but rather a shift in the deducted amount between partners. This is clear in Figure A1 in the Appendix. Figure A1a shows the take-up rate of the declared deduction by primary earners at kink 1, which includes any lifelong learning expenditures that are shifted from the secondary earner to the primary earner. There is a clear upward jump in the take-up rate. Also for kink 2, we observe a significant upward jump in the take-up rate, see Figure A1b.[27] Taken at face value, this would suggest a very large positive effect on the take-up of lifelong learning of the tax deduction. However, for secondary earners we then observe a counterintuitive decline in the take-up rate of the deduction for lifelong learning expenditures to the right of kink, both for kink 1 and 2, see Figure A1c and A1d, respectively. These estimates also include shifting from the secondary earner to the primary earner. Therefore, we consider the effect on the *own* (declared) lifelong learning expenditures next.[28]

Figure 5.3c and 5.3d display much smaller jumps in the take-up rate of *own* lifelong learning expenditures of primary earners at kink 1 and 2 than for the take-up rate of *declared* lifelong learning expenditures. Figure 5.3c shows that the treatment effect is +0.42 percentage points (+23.3%), and figure 5.3d shows that for kink 2 the treatment effect is +0.18 percentage points (+10%) using a donut hole of 1,000 euro.[29] In our preferred specification including controls, the effect remains statistically significant and similar in size (see Table A1 in the appendix).[30]

Figure 5.3e and 5.3f show the effect on the average own amount deducted for primary earners around kink 1 and 2, respectively. These figures also suggest a much smaller effect than on the declared amount (compare with Figure A2 in the appendix), but still positive and statistically significant. The treatment effect for kink 1 is +6.3 euro (+28.8%) and for kink 2 it is +7.9 euro (+20.2%).[31] Regression results are given in Table A1 in the appendix and confirm the graphical analysis.

[27]There is also a clear upward jump in the average amount deducted for kink 1 and 2, see Figure A2.

[28]We still need to apply a donut-RD design though, because manipulation of the income variable still affects the composition of primary earners close to the kink.

[29]There is no apparent excess mass beyond 1,000 euro away from the kink.

[30]The estimates in Table A1 can also be converted into an elasticity of (declared own) lifelong learning expenditures with respect to the effective costs of lifelong learning expenditures. At kink 2, the effective costs are 11% lower on the right of the kink than on the left of the kink. The effective costs on the left-hand side are $(1 - 0.42) * (2,225 - 500) + 500 = 1,500$ euro, where 2,225 are the average own lifelong learning expenditures around kink 2. The effective costs on the right-hand side are $(1 - 0.52) * (2,225 - 500) + 500 = 1,328$ euro. This suggests an elasticity at kink 2 of $+7.9/(-11) \approx -0.7$.

[31]At kink 2, the corresponding elasticity of (declared own) lifelong learning with respect to the effective costs (based on the preferred regression specification) $+16.4/(-11) \approx -1.5$.

The appendix also reports the figures and table of the treatment effects on the own lifelong learning expenditures of secondary earners (see Figure A3 and Table A2). For both kinks we find a very small and statistically insignificant effect for the own deduction, as opposed to the statistically significant negative effects for the declared deduction.

In the end, although we think the results for couples are interesting and the results for the own deductions are plausible in terms of the sign of the effect, we have the most confidence in the estimates for singles when it comes to the causal effect of the tax deduction on lifelong learning expenditures, because we do not observe manipulation of the running variable for singles.

5.6.3 Discussion

An important question that remains is whether we measure the effect on actual lifelong learning expenditures, or simply the reporting of lifelong learning expenditures. Indeed, the financial incentive to file expenditures is higher for people above the kink than for people below the kink. To investigate this question further, we linked survey data from the Labour Force Survey to the tax return data. Specifically, the Labour Force Survey contains self-reported participation in post-initial training and whether or not the person had out-of-pocket expenses for post-initial training.

Using our preferred specifications we find suggestive evidence at both kink 1 and kink 2 of an increase in self-paid training, but no change in employer-paid training, see Table 5.8. Because we have to use survey data for this analysis, the number of observations is much smaller and the estimates are not very robust. Nevertheless, the analysis does not reject the hypothesis that we measure actual increases in lifelong learning expenditures. Furthermore, the estimated treatment effects by age groups presented in Table 5.7 are also consistent with differences in actual spending on lifelong learning expenditures. Middle-aged persons, for whom the returns to lifelong learning expenditures are relatively large (compared to e.g. older workers), are more likely to have a difference in lifelong learning expenditures than for example older persons, whereas the difference in financial incentives to the left- and right-hand side of the kink are the same for both groups.

An additional piece of evidence that people using the deductible actually do take up more training is that enrollment in publicly funded education is substantially higher for people using the deductible than for those who do not. Table 5.9 reports enrollment rates for people using the deductible and those not declaring any training expenditures. We find that from post-secondary education (ISCED level 4) upwards enrollment rates are significantly higher for those using the deductible. The remaining 78% are either not enrolled or enrolled in private education, for which we have no data.

Finally, people also predominantly report using the deductible for training expenses. Table 5.10 shows the frequency of voluntarily given descriptions accompanying a random sample of 50,000 deductions in 2013. We find that about 23% report using the deductible for tuition fees or similar terms, 14% for books and 9% for "study costs". People also often report the names of private education institutes or universities.

Table 5.8: Treatment effect estimates for singles on the probability to participate in and pay for training: Labour Force Survey

	(1) Training	(2) Partly self-paid training	(3) Self-paid training	Observations
Panel A. Kink 1				
Bandwidth				
1,000	−0.0051	0.0214	0.0182	5,816
	(0.0317)	(0.0141)	(0.0131)	
1,330	−0.0085	0.0226**	0.0241***	7,732
	(0.0245)	(0.0114)	(0.0091)	
1,500	−0.0072	0.0237***	0.0226***	8,756
	(0.0175)	(0.0081)	(0.0064)	
Panel B. Kink 2				
Bandwidth				
1,500	−0.0537	0.0073	0.0057	1,868
	(0.0522)	(0.0261)	(0.0251)	
2,000	0.0090	0.0257*	0.0227*	2,468
	(0.0360)	(0.0151)	(0.0137)	
2,500	−0.0266	0.0098	0.0016	3,078
	(0.0255)	(0.0123)	(0.0111)	

Notes: Sample period 2006–2012. Cluster-robust standard errors clustered by income bins of 100 euro in parentheses, *** $p<0.01$, ** $p<0.05$, * $p<0.1$. The data consists of our sample merged to data from the Labour Force Survey. We use two questions to determine training status: whether workers currently take part in training, or whether they have done so in the past four weeks. We then use information on who pays for the training to determine our outcome variable. All regressions include year fixed effects and controls for gender, ethnicity, age, age^2 and the number of children in the household.

Another question is why the effect seems to be smaller for low-income singles than for high-income singles. One potential explanation is that differences in statutory tax rates are perhaps less important than differences in effective marginal tax rates due to income-dependent tax credits and subsidies for low-income singles. Indeed, Figure A11a and A11b in the appendix, for childless singles and lone parents respectively, show that effective marginal tax rates are quite different from statutory tax rates for low-income singles (for high-income singles they are almost identical in the period

Table 5.9: Enrollment in publicly funded education (%)

	Using the deductible	Not using the deductible
ISCED 2011 levels		
2 (Lower secondary)	0.03	0.08
3 (Upper secondary)	0.12	0.27
4 (Post-sec, non-tertiary)	2.59	1.29
5 (Short tertiary)	14.32	0.50
6 (Bachelor)	4.78	0.22
7 (Master)	0.13	0.00
Total	21.97	2.37
Not enrolled or enrolled in private education	78.03	97.63
N	831,429	34,367,021

Notes: Data for 2006–2012, all people between 25 and 60 years old, excluding people eligible for student aid. Enrollment rates are calculated from administrative records available through Statistics Netherlands.

Table 5.10: Frequency of words reported in tax filings

Keyword	% reported
Tuition fees (& similar terms)	23
Books	14
Study costs	9
Several private education institutes	3
University of applied science	3
Study materials	2
Computer	2
University	2
Master	1
Management	1
Academy	1
Dutch	1

Notes: Calculated from a random sample of 50,000 voluntarily entered descriptions in the tax filings of 2013 made available by the Ministry of Finance.

we consider). As a result, low-income singles may not respond to the differences in *statutory* marginal tax rates.

One could also argue that the tax credit for lifelong learning is not very salient for low-income singles or that friction costs prevent them from filing lifelong learning expenditures (Ladner et al., 2009; Chetty et al., 2013). One piece of evidence that the tax credit is not very salient, is that of the full population in the Labor Force Survey who claim that they took (partly) self-paid training during the year, and who hence are potentially eligible for the tax credit, only about a quarter of them actually declare their expenditures. This is similar for low- and high-income individuals. In addition, when we look at the change in the system in 2013, when the threshold was reduced from 500 euro (in 2012) to 250 euro (2013), we see a sharp increase in filed lifelong learning expenditures between 250 and 500 euro, both for low- and high-income singles, see Figure A12. This suggests that the tax deduction was as salient to low-income singles as it was to high-income singles. However, the difference in effective marginal tax rates may have been less salient.

Filing frictions also appear to be small for both low- and high-income groups. We observe that both groups also report very small expenditures on lifelong learning expenditures, see Figure A12 in the appendix. Another reason why low-income individuals may be less likely to respond has to do with other costs associated with post-initial education. Time constraints may be considerable, especially when young children are present, and many people - particularly lower-educated individuals - dislike formal learning and hence may experience psychic costs from doing so. Furthermore, people may be myopic – again particularly lower-educated individuals – or more generally underestimate the gains of lifelong learning (see e.g. Heckman et al., 2006).

5.7 Conclusion

In this paper we have studied the effectiveness of a tax deduction for lifelong learning expenditures in terms of the take-up rate of lifelong learning expenditures and the average amount of lifelong learning expenditures. For singles, which is our preferred group because for them we do not observe manipulation of the running variable, we find heterogeneous effects of the tax deduction. For the high-income group the take-up rate of lifelong learning expenditures increases by 10%. The effect is particularly large for men, natives, higher educated and middle-aged. However, at a relatively low level of income the additional effect of the tax deduction is essentially zero, with tight confidence intervals.

For couples, we find large positive treatment effects for the *declared* deduction for primary earners, and counterintuitive negative treatment effects for secondary earners. However, this is due to shifting of the lifelong learning expenditures from secondary earners to primary earners. Indeed, when we consider the take-up rate of *own* lifelong learning expenditures instead, we find smaller effects for primary earners and a negligible effect for secondary earners. However, manipulation of the running variable in couples is still problematic, and our ‘solution’ of applying donut-RD regressions, where we leave out the bins with excess mass close to the discontinuity, has the downside of comparing groups to the left and the right of the discontinuity that are increasingly dissimilar. Hence, we prefer the estimates for singles to learn about the causal effect of the tax deduction on lifelong learning expenditures.

Our study contributes to the scarce literature on tax incentives for lifelong learning. We show that, at the margin, tax incentives provide an incentive for high-income workers to pursue training. For low-income workers tax incentives do not increase their level of training. Overall we might wonder whether our findings show that the policy is effective. On the one hand, the marginal deadweight loss seems rather high. At kink 1 it is 100%, since we do not find an increase in training at kink 1. At kink 2 it is around 90% at 1,000 euros above the kink. Compared to the literature on schooling vouchers, which shows deadweight losses of 30% (Schwerdt et al., 2012) to 60% (Hidalgo et al., 2014), the fiscal incentive seems less effective. The apparent lack of effectiveness is in line with the literature on the effects of fiscal incentives for initial education (Bulman and Hoxby, 2015). However, the deadweight loss does not take into account the size of the effect relative to the costs. Indeed, the elasticity of lifelong learning with respect to effective costs is non-negligible. At kink 2, effective costs of life long learning drop by 6%, while take-up of the deductible increases by 10%. This implies a (sizeable) elasticity of −1.7 for higher incomes.[32]

[32] For a full cost-benefit analysis we would need to know the return on these investments. We have used our estimated relation between training expenses and distance to the kink as a ‘first stage’ to estimate the subsequent effect on income growth. However, we find no significant effects on income growth.

A Appendix

Figure A1: Declared deductible for primary and secondary earners

Kink 1 | Kink 2

Declared deductible primary earners

Estimate of the discontinuity: 0.0078 (0.0016)***

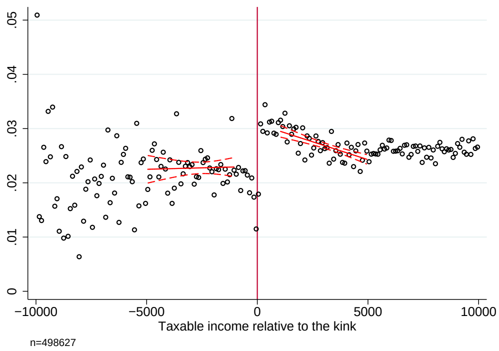

(a)

Estimate of the discontinuity: 0.0081 (0.0012)***

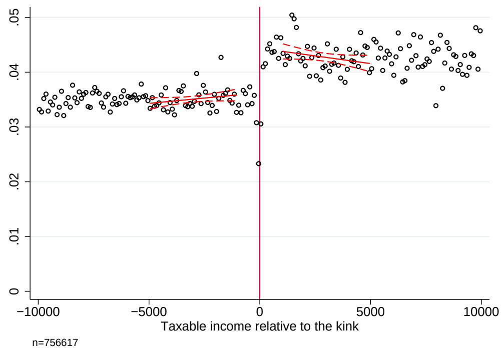

(b)

Declared deductible secondary earners

Estimate of the discontinuity: −0.0042 (0.0012)***

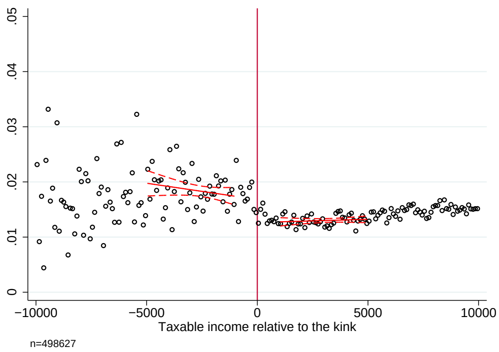

(c)

Estimate of the discontinuity: −0.0052 (0.0007)***

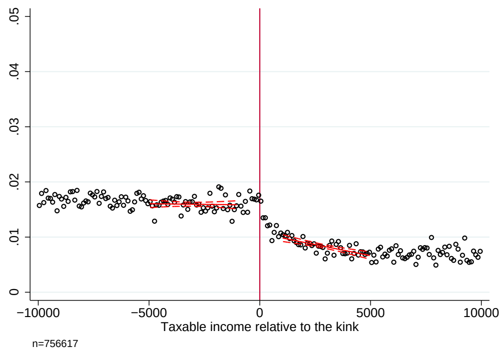

(d)

Notes: Own calculations based on register data from Statistics Netherlands. The regression lines are linear functions without any control variables, with a separate intercept and slope on the right-hand side of the kink and using a donut hole of 1000 euro on either side of the kink. The deducted amount includes the zeros for non-users. The estimate of the discontinuity without demographic control variables is presented above each figure.

Table A1: Treatment effect estimates for primary earners on the probability to use the deductible and the deducted amount, for different widths of the donut hole

	Without controls			With controls			Observations
	(1)	(2)	(3)	(4)	(5)	(6)	
	Linear	Quadratic	Local linear	Linear	Quadratic	Local linear	
Kink 1							
Panel A: Take-up rate of the deduction							
Donut 1500 euro	0.0041***	0.0040**	0.0037**	0.0032**	0.0040**	0.0037**	444, 991
	(0.0014)	(0.0018)	(0.0018)	(0.0014)	(0.0017)	(0.0017)	
Donut 1000 euro	0.0049***	0.0045***	0.0042***	0.0042***	0.0046***	0.0042***	498, 627
	(0.0012)	(0.0015)	(0.0015)	(0.0012)	(0.0015)	(0.0015)	
Donut 500 euro	0.0042***	0.0038***	0.0038***	0.0038***	0.0041***	0.0039***	553, 066
	(0.0009)	(0.0010)	(0.0009)	(0.0009)	(0.0010)	(0.0009)	
Donut 0 euro	0.0050***	0.0053***	0.0053***	0.0051***	0.0057***	0.0057***	624, 570
	(0.0013)	(0.0015)	(0.0014)	(0.0015)	(0.0016)	(0.0015)	
Panel B: Deducted amount (in euro)							
Donut 1500 euro	5.4107**	5.2795*	5.5700**	4.1866**	5.3381*	5.5644**	444, 991
	(2.1359)	(2.8945)	(2.8332)	(2.1575)	(2.9072)	(2.8542)	
Donut 1000 euro	5.6843***	6.0792***	6.2702***	4.8075***	6.2164***	6.3304***	498, 627
	(1.8753)	(2.8135)	(2.8308)	(1.8799)	(2.8115)	(2.8373)	
Donut 500 euro	4.4841***	4.5806***	4.7594***	3.9581***	4.8480***	4.9424***	553, 066
	(1.3746)	(1.7221)	(1.6641)	(1.3673)	(1.7095)	(1.6607)	
Donut 0 euro	4.9564***	5.7033***	5.9056***	4.9354**	6.2277***	6.3219***	624, 570
	(2.8674)	(3.1811)	(2.7430)	(3.1465)	(3.3984)	(2.8847)	
Kink 2							
Panel C: Take-up rate of the deduction							
Donut 1500 euro	−0.0002	0.0001	0.0002	0.0011	0.0013	0.0014	660, 928
	(0.0016)	(0.0016)	(0.0016)	(0.0016)	(0.0016)	(0.0015)	
Donut 1000 euro	0.0014	0.0018*	0.0018*	0.0026**	0.0029***	0.0030***	756, 617
	(0.0011)	(0.0011)	(0.0010)	(0.0011)	(0.0010)	(0.0010)	
Donut 500 euro	0.0019**	0.0022***	0.0022***	0.0032***	0.0034***	0.0034***	854, 511
	(0.0007)	(0.0007)	(0.0007)	(0.0007)	(0.0007)	(0.0007)	
Donut 0 euro	0.0023**	0.0088***	0.0023**	0.0038**	0.0037**	0.0036**	970, 301
	(0.0011)	(0.0011)	(0.0011)	(0.0016)	(0.0016)	(0.0015)	
Panel D: Deducted amount (in euro)							
Donut 1500 euro	4.1094	4.8658	4.9106	6.6603*	7.3894**	7.4165**	660, 928
	(3.5941)	(3.5063)	(3.4895)	(3.5560)	(3.4786)	(3.4682)	
Donut 1000 euro	7.2399***	7.9263***	7.9087***	9.6271***	10.2659***	10.2346***	756, 617
	(2.7686)	(2.7847)	(2.8005)	(2.7557)	(2.7823)	(2.8003)	
Donut 500 euro	6.3917***	6.7515***	6.6859***	8.8987***	9.1924***	9.1108***	854, 511
	(1.9854)	(2.0319)	(2.0471)	(1.9828)	(2.0335)	(2.0461)	
Donut 0 euro	6.3236**	6.1236**	5.8874**	9.2013***	8.7722***	8.4473***	970, 301
	(2.4934)	(2.4747)	(2.3347)	(3.3595)	(3.2780)	(3.0388)	

Notes: Sample period 2006–2012. Cluster-robust standard errors clustered by income bins of 100 euro in parentheses, *** p<0.01, ** p<0.05, * p<0.1. All regressions include year fixed effects. Columns (1)-(3) are without demographic control variables, columns (4)-(6) are with demographic control variables. Columns (1) and (4) assume a linear relation between taxable income and the dependent variable and allow for a discontinuity in the intercept. Columns (2) and (5) assume a quadratic relation between taxable income and the dependent variable and allow for a discontinuity in the intercept. Columns (3) and (6) assume a linear relation between taxable income and the dependent variable and allow for a discontinuity in the intercept and in the linear relation between taxable income and the dependent variable.

Table A2: Treatment effect estimates secondary earners on the probability to use the deductible and the deducted amount, for different widths of the donut hole

	Without controls			With controls			Obs.
	(1) Linear	(2) Quadratic	(3) Local linear	(4) Linear	(5) Quadratic	(6) Local linear	
Kink 1							
Panel A: Take-up rate of the deduction							
Donut 1500 euro	−0.0008	0.0012	0.0012	−0.0018	0.0013	0.0012	444,991
	(0.0011)	(0.0017)	(0.0017)	(0.0012)	(0.0017)	(0.0017)	
Donut 1000 euro	−0.0003	0.0010	0.0007	−0.0008	0.0012	0.0009	498,627
	(0.0010)	(0.0013)	(0.0013)	(0.0010)	(0.0014)	(0.0013)	
Donut 500 euro	0.0000	0.0008	0.0006	−0.0003	0.0011	0.0009	553,066
	(0.0008)	(0.0010)	(0.0010)	(0.0008)	(0.0010)	(0.0010)	
Donut 0 euro	0.0005	0.0013	0.0013	0.0007	0.0020	0.0019	624,570
	(0.0013)	(0.0014)	(0.0012)	(0.0016)	(0.0016)	(0.0014)	
Panel B: Deducted amount (in euro)							
Donut 1500 euro	−2.6145	−0.9536	−1.0468	−3.8541*	−0.8615	−1.0158	444,991
	(2.2672)	(3.3447)	(3.2492)	(2.2766)	(3.3831)	(3.2931)	
Donut 1000 euro	−2.1589	−0.3692	−0.3626	−2.9122	0.0100	−0.0662	498,627
	(1.7964)	(2.4282)	(2.3334)	(1.8436)	(2.4664)	(2.3683)	
Donut 500 euro	0.2842	2.0511	1.8645	−0.0813	2.5715	2.2934	553,066
	(1.5644)	(1.8637)	(1.7838)	(1.6450)	(1.9031)	(1.8161)	
Donut 0 euro	1.1794	2.7829	2.6597	1.5290	3.8269	3.5936*	624,570
	(2.0516)	(2.1262)	(1.8443)	(2.4663)	(2.4987)	(2.1267)	
Kink 2							
Panel C: Take-up rate of the deduction							
Donut 1500 euro	−0.0016	−0.0008	−0.0008	−0.0006	0.0001	0.0001	660,928
	(0.0018)	(0.0016)	(0.0016)	(0.0017)	(0.0016)	(0.0015)	
Donut 1000 euro	−0.0001	0.0003	0.0002	0.0008	0.0012	0.0011	756,617
	(0.0012)	(0.0012)	(0.0012)	(0.0012)	(0.0012)	(0.0012)	
Donut 500 euro	0.0005	0.0008	0.0008	0.0015	0.0018*	0.0017*	854,511
	(0.0010)	(0.0009)	(0.0009)	(0.0010)	(0.0009)	(0.0009)	
Donut 0 euro	0.0010	0.0012	0.0012	0.0021*	0.0022*	0.0022*	970,301
	(0.0009)	(0.0009)	(0.0009)	(0.0012)	(0.0012)	(0.0012)	
Panel D: Deducted amount (in euro)							
Donut 1500 euro	−4.6525	−4.0725	−4.0050	−2.7501	−2.1733	−2.1232	660,928
	(3.2716)	(3.5188)	(3.4989)	(3.2642)	(3.5187)	(3.5050)	
Donut 1000 euro	−1.6134	−1.2323	−1.2399	0.1965	0.5499	0.5273	756,617
	(2.2687)	(2.3877)	(2.3788)	(2.2757)	(2.4061)	(2.4001)	
Donut 500 euro	0.3518	0.3819	0.2829	2.2895	2.2731	2.1584	854,511
	(1.7042)	(1.7657)	(1.7439)	(1.7508)	(1.8065)	(1.7786)	
Donut 0 euro	1.8676	1.9498	1.8870	4.0630*	3.9716*	3.8371*	970,301
	(1.6805)	(1.6670)	(1.6318)	(2.1608)	(2.0930)	(1.9981)	

Notes: Sample period 2006–2012. Cluster-robust standard errors clustered by income bins of 100 euro in parentheses, *** $p<0.01$, ** $p<0.05$, * $p<0.1$. All regressions include year fixed effects. Columns (1)-(3) are without demographic control variables, columns (4)-(6) are with demographic control variables. Columns (1) and (4) assume a linear relation between taxable income and the dependent variable and allow for a discontinuity in the intercept. Columns (2) and (5) assume a quadratic relation between taxable income and the dependent variable and allow for a discontinuity in the intercept. Columns (3) and (6) assume a linear relation between taxable income and the dependent variable and allow for a discontinuity in the intercept and in the linear relation between taxable income and the dependent variable.

Table A3: Full estimation results for the preferred specification for singles

	(1)	(2)	(3)	(4)
	Kink 1 (bandwidth 1,330)		Kink 2 (bandwidth 2,000)	
	Deductible	Deducted amount	Deductible	Deducted amount
Above the kink x taxable income	-0.0014	-2.4376	0.0038^{***}	5.8728
	(0.0012)	(2.3543)	(0.0010)	(3.8610)
Taxable income	0.0005	0.0336	-0.0011^{*}	-1.6255
	(0.0007)	(1.2078)	(0.0006)	(2.3984)
Controls				
Age	-0.0062^{***}	-9.2879^{***}	-0.0028^{***}	-19.7179^{***}
	(0.0003)	(0.4486)	(0.0004)	(3.7302)
Age^2	0.0001^{***}	0.0816^{***}	0.0000^{***}	0.1760^{***}
	(0.0000)	(0.0050)	(0.0000)	(0.0402)
Female	0.0037^{***}	-1.1878	0.0222^{***}	34.9918^{***}
	(0.0008)	(1.4376)	(0.0010)	(3.6599)
Foreign	0.0016^{***}	2.6464^{**}	0.0045^{***}	24.0090^{***}
	(0.0006)	(1.0514)	(0.0011)	(4.9127)
Number of children	-0.0035^{***}	-5.8402^{***}	-0.0021^{***}	0.5144
	(0.0003)	(0.5918)	(0.0004)	(2.2350)
Constant	0.1846^{***}	271.3887^{***}	0.1175^{***}	556.1877^{***}
	(0.0065)	(10.0242)	(0.0097)	(83.2903)
Observations	662,848	662,848	197,584	197,584

Notes: Sample period 2006–2012. Cluster-robust standard errors clustered by income bins of 100 euro in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Year fixed effects included.

Table A4: Treatment effect estimates for singles: standard errors ‘clustered’ at the individual level

	(1)	(2)	(3)	(4)	
	Use of the deductible		Deducted amount (in euro)		
	No controls	Controls	No controls	Controls	Observations
Panel A. Kink 1					
Bandwidth					
1,000	0.0006	0.0007	-2.2945	-2.3580	496,957
	(0.0017)	(0.0016)	(3.4489)	(3.4383)	
1,330	-0.0014	-0.0014	-2.4042	-2.4376	662,848
	(0.0011)	(0.0011)	(2.2193)	(2.2130)	
1,500	-0.0006	-0.0006	-1.4232	-1.5225	749,526
	(0.0009)	(0.0009)	(1.8398)	(1.8341)	
Panel B. Kink 2					
Bandwidth					
1,500	0.0024	0.0024	4.5190	4.8451	148,526
	(0.0023)	(0.0023)	(9.2539)	(9.2401)	
2,000	0.0038^{**}	0.0038^{**}	5.5721	5.8728	197,584
	(0.0015)	(0.0015)	(6.2202)	(6.1994)	
2,500	0.0031^{***}	0.0032^{***}	7.8314^{*}	8.0554^{*}	246,949
	(0.0011)	(0.0011)	(4.6452)	(4.6343)	

Notes: Sample period 2006–2012. Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. All regressions include year fixed effects. The regressions with controls include gender, ethnicity, age, age^2 and the number of children in the household as demographic controls.

Figure A2: Declared deducted amount for primary and secondary earners

Kink 1 | Kink 2

Primary earners

Estimate of the discontinuity: 13.5425 (2.7925)***

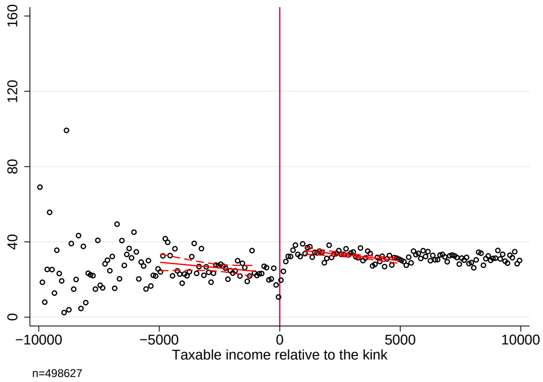

(a)

Estimate of the discontinuity: 15.3107 (3.2629)***

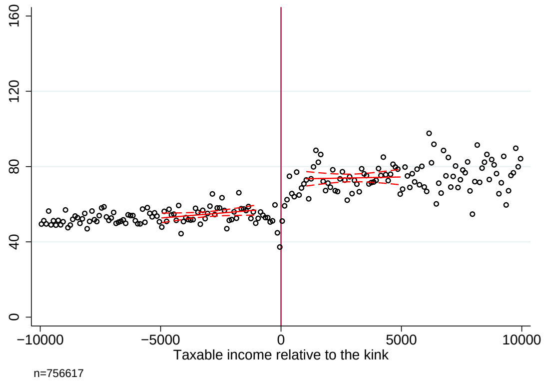

(b)

Secondary earners

Estimate of the discontinuity: -7.4534 (1.7046)***

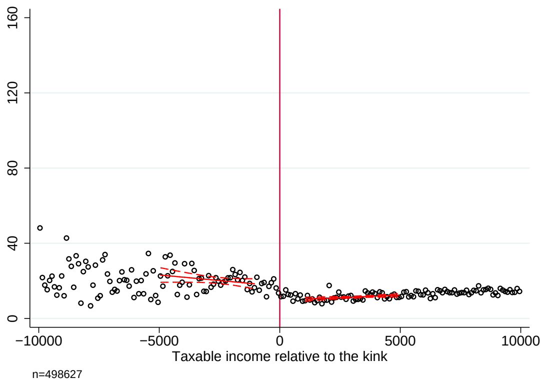

(c)

Estimate of the discontinuity: -8.7024 (1.4860)***

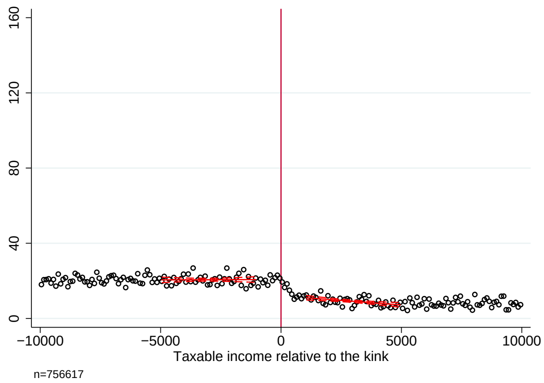

(d)

Notes: Own calculations based on register data from Statistics Netherlands. The regression lines are linear functions without any control variables, with a separate intercept and slope on the right-hand side of the kink and using a donut hole of 1000 euro on either side of the kink. The deducted amount includes the zeros for non-users. The estimate of the discontinuity without demographic control variables is presented above each figure.

Figure A3: Own use of the deductible and own amount for secondary earners

Kink 1 | Kink 2

Own deductible

Estimate of the discontinuity: 0.0008 (0.0014)

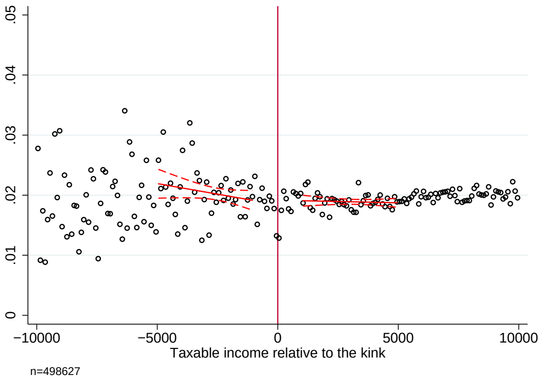

(a)

Estimate of the discontinuity: 0.0003 (0.0010)

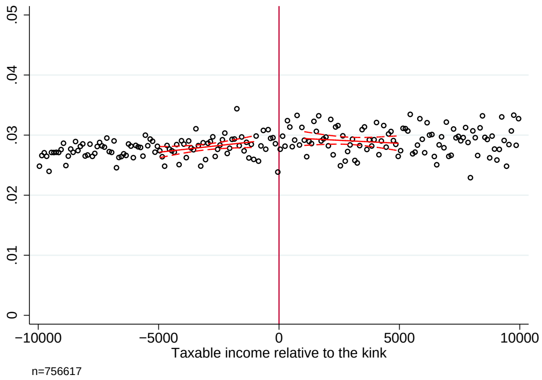

(b)

Own deducted amount

Estimate of the discontinuity: −.2612 (2.360)

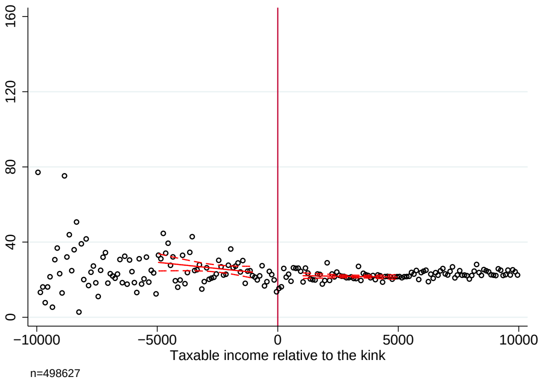

(c)

Estimate of the discontinuity: −1.3002 (2.3858)

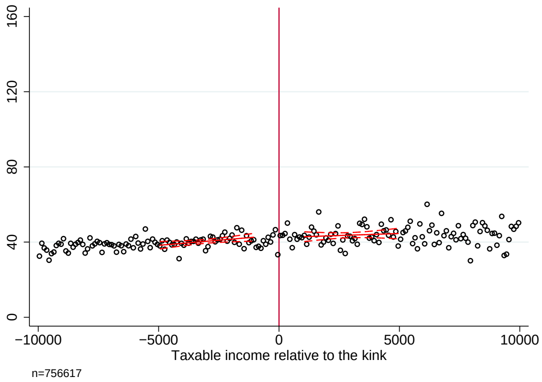

(d)

Notes: Own calculations based on register data from Statistics Netherlands. The regression lines are linear functions without any control variables, with a separate intercept and slope on the right-hand side of the kink and using a donut hole of 1000 euro on either side of the kink. The deducted amount includes the zeros for non-users. The estimate of the discontinuity without demographic control variables is presented above each figure.

Figure A4: Average (statutory) marginal tax rates in subsequent years for the sample around the kink in 2006

Kink 1 Kink 2

Singles

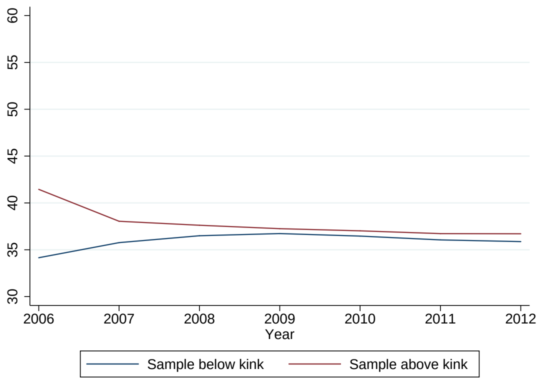

(a) Singles around kink 1

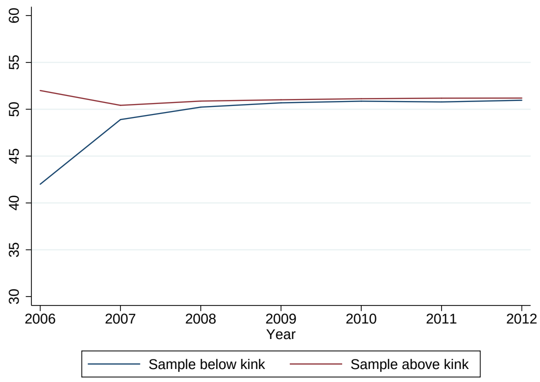

(b) Singles around kink 2

Primary earners

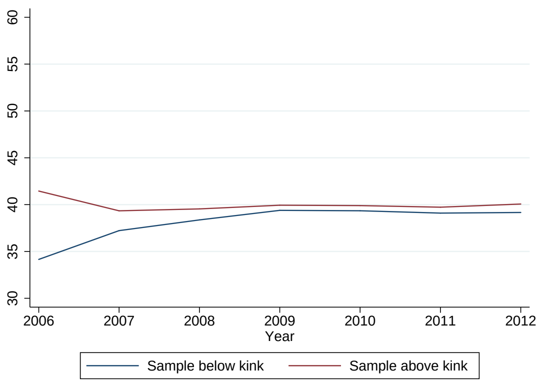

(c) Primary earners around kink 1

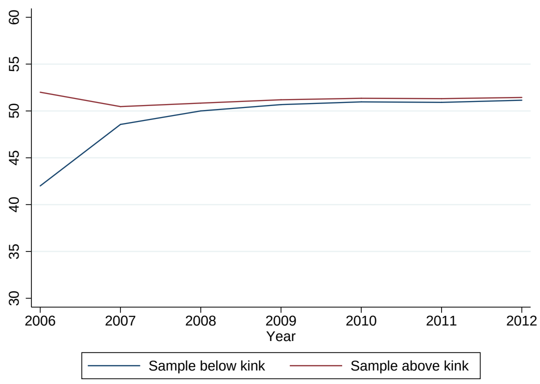

(d) Primary earners around kink 2

Notes: Own calculations based on tax return data from Statistics Netherlands.

Figure A5: RKD plots for control variables for singles

Kink 1 — Kink 2

Female

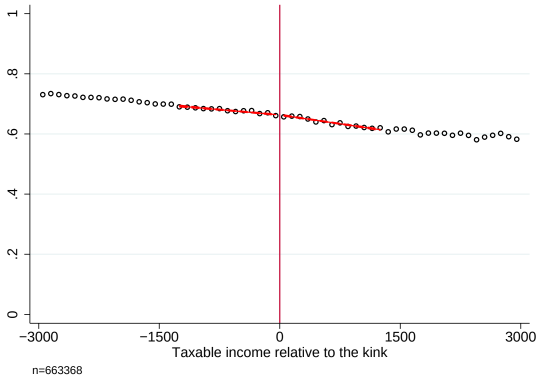

(a)

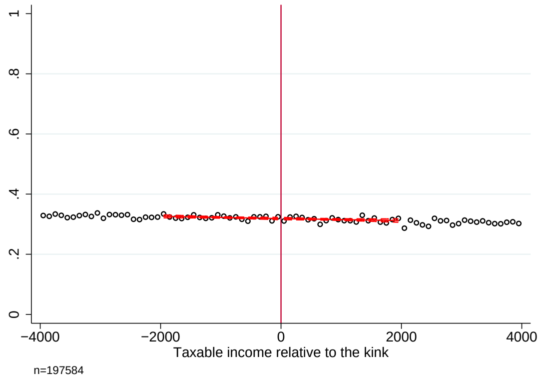

(b)

Age

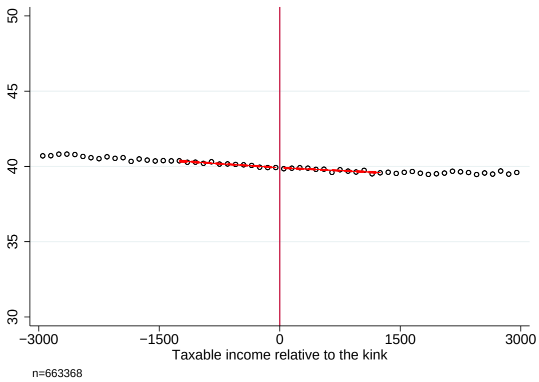

(c)

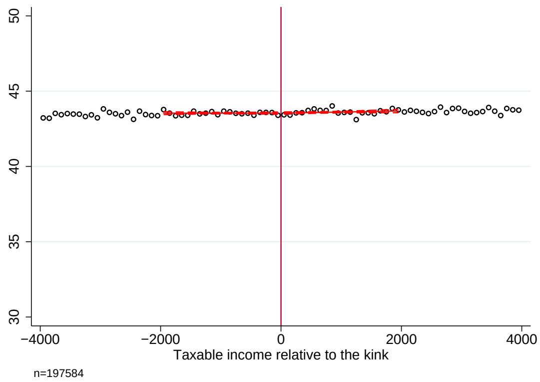

(d)

Notes: Own calculations based on register data from Statistics Netherlands. The regression lines are linear functions without any control variables, with a separate intercept and slope on the right-hand side of the kink. The deducted amount includes the zeros for non-users. Estimates for kink 1 include observations from minus 1,330 to plus 1,330 euro relative to the kink. Estimates for kink 2 include observations from minus 2,000 to plus 2,000 euro relative to the kink.

Figure A5: RKD plots for control variables for singles (cont.)

Kink 1 | Kink 2

Foreign

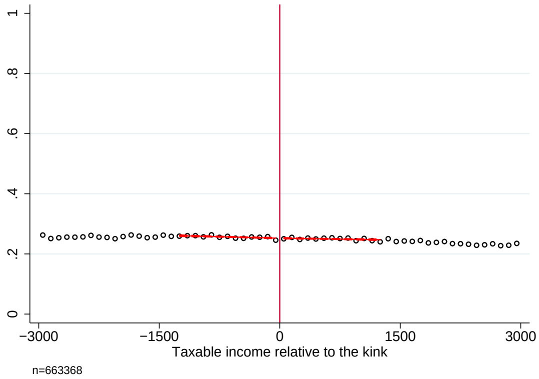

(e)

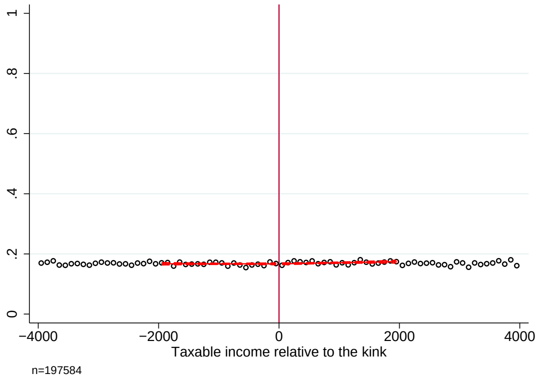

(f)

Number of children

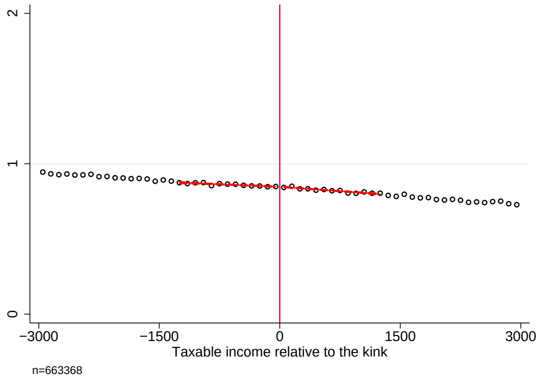

(g)

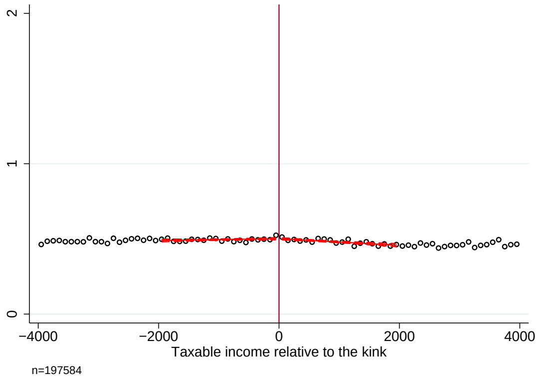

(h)

Notes: Own calculations based on register data from Statistics Netherlands. The regression lines are linear functions without any control variables, with a separate intercept and slope on the right-hand side of the kink. The deducted amount includes the zeros for non-users. Estimates for kink 1 include observations from minus 1,330 to plus 1,330 euro relative to the kink. Estimates for kink 2 include observations from minus 2,000 to plus 2,000 euro relative to the kink.

Figure A6: RD plots for control variables for the primary earner in a couple

Kink 1 Kink 2

Female

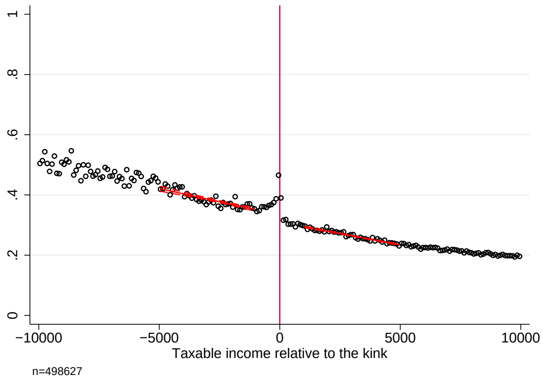

(a)

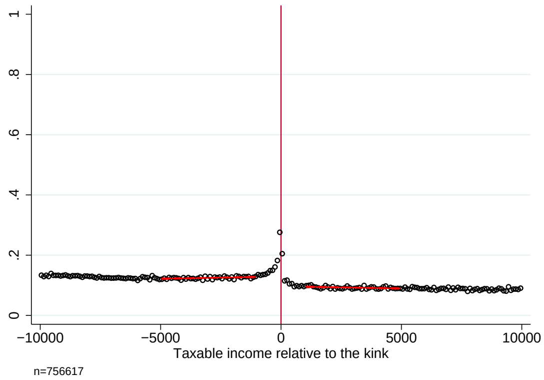

(b)

Age

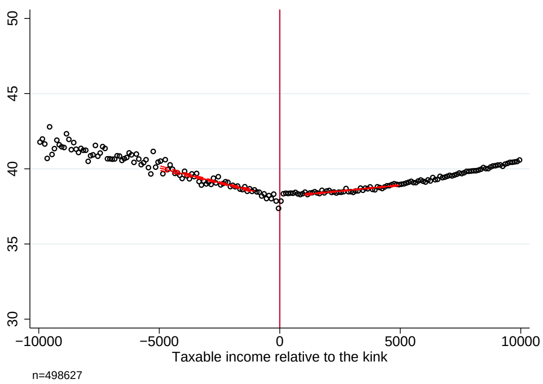

(c)

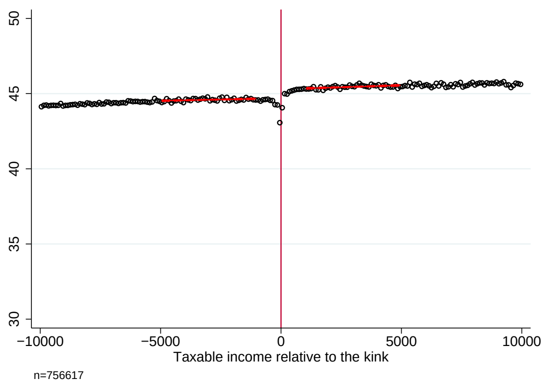

(d)

Notes: Own calculations based on register data from Statistics Netherlands. The regression lines are linear functions without any control variables, with a separate intercept and slope on the right-hand side of the kink and using a donut hole of 1000 euro on either side of the kink.

Figure A6: RD plots for control variables for the primary earner in a couple (cont.)

Kink 1 Kink 2

Foreign

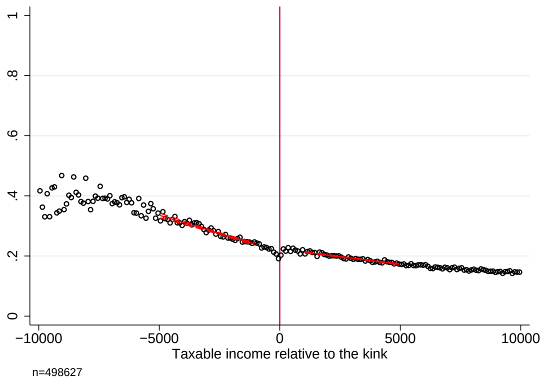

(e)

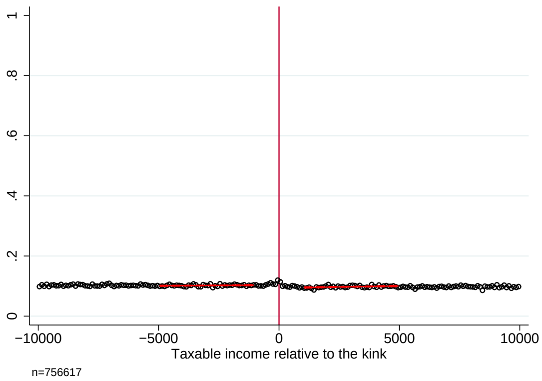

(f)

Number of children

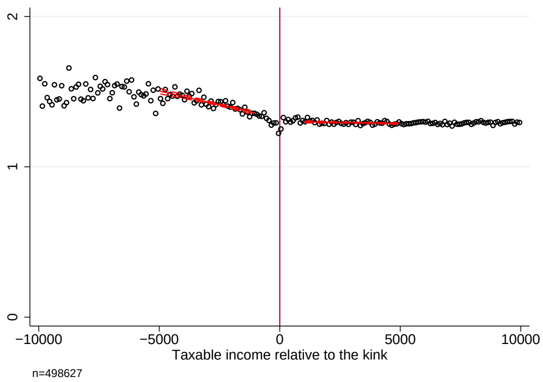

(g)

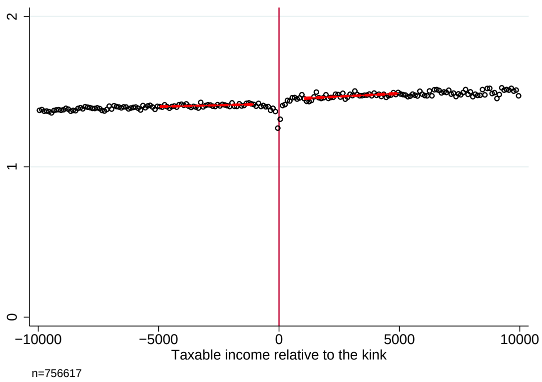

(h)

Notes: Own calculations based on register data from Statistics Netherlands. The regression lines are linear functions without any control variables, with a separate intercept and slope on the right-hand side of the kink and using a donut hole of 1000 euro on either side of the kink.

Figure A7: RD plots for control variables for the secondary earner in a couple

Kink 1 | Kink 2

Female

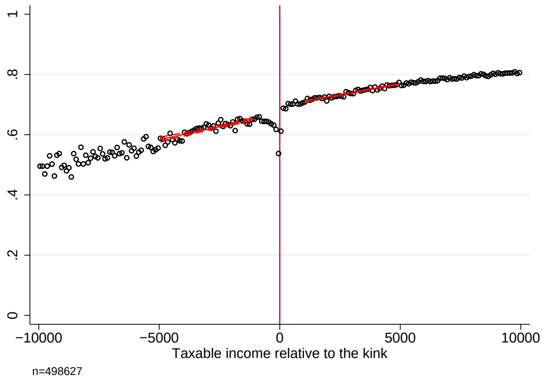

(a)

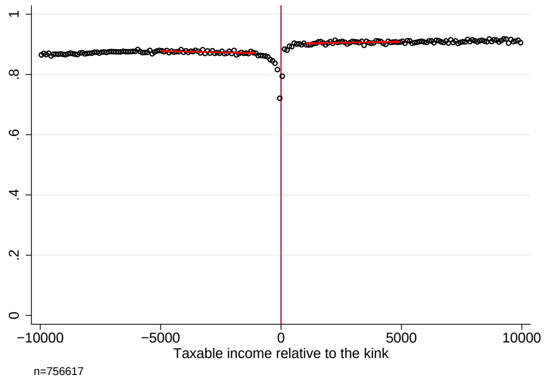

(b)

Age

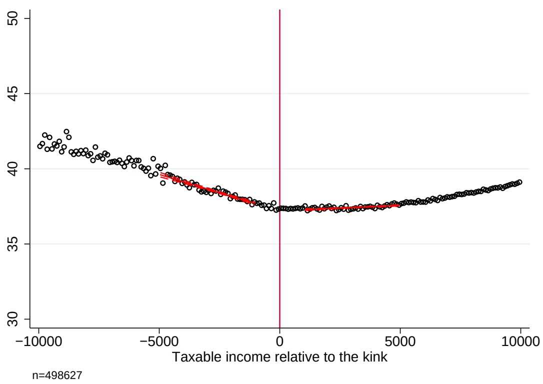

(c)

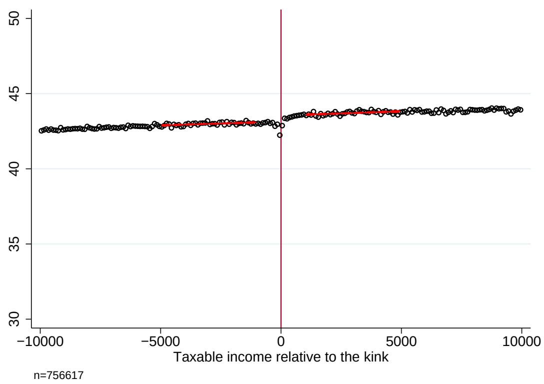

(d)

Notes: Own calculations based on register data from Statistics Netherlands. The regression lines are linear functions without any control variables, with a separate intercept and slope on the right-hand side of the kink and using a donut hole of 1000 euro on either side of the kink.

Figure A7: RD plots for control variables for the secondary earner in a couple (cont.)

Kink 1 Kink 2

Foreign

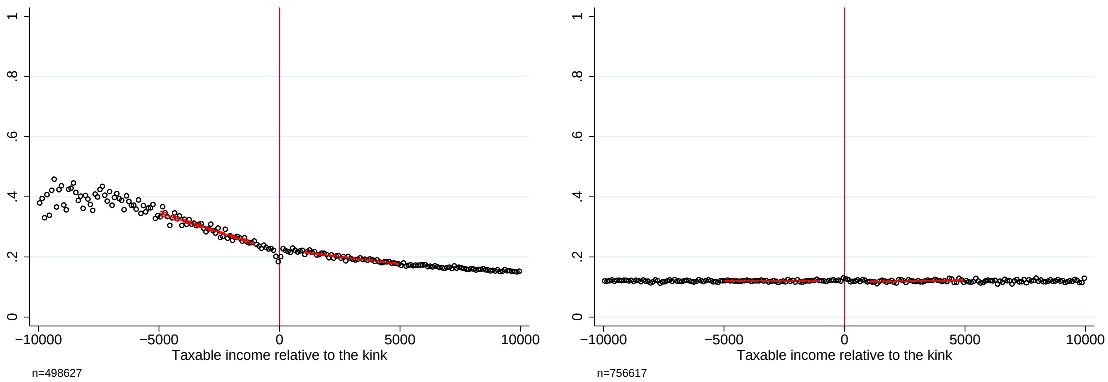

(e) (f)

Number of children

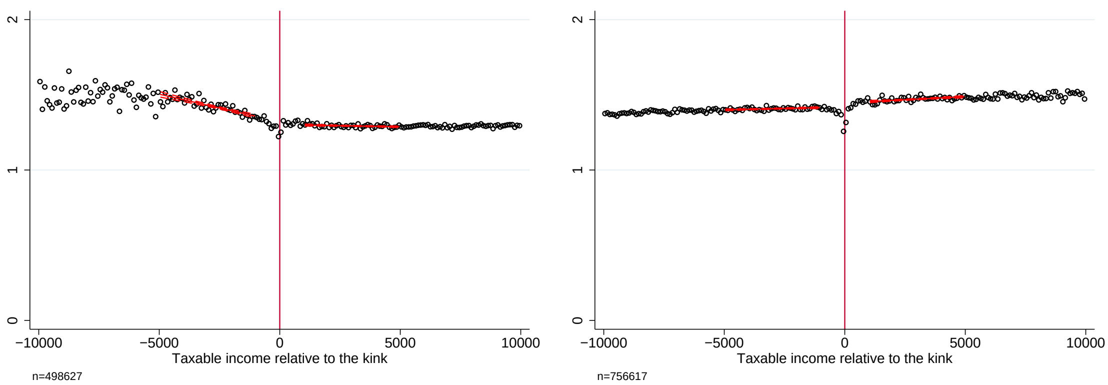

(g) (h)

Notes: Own calculations based on register data from Statistics Netherlands. The regression lines are linear functions without any control variables, with a separate intercept and slope on the right-hand side of the kink and using a donut hole of 1000 euro on either side of the kink. The estimate of the discontinuity without demographic control variables is presented above each figure.

Figure A8: Average deducted amount for those who take up the deduction

Kink 1 Kink 2

Singles

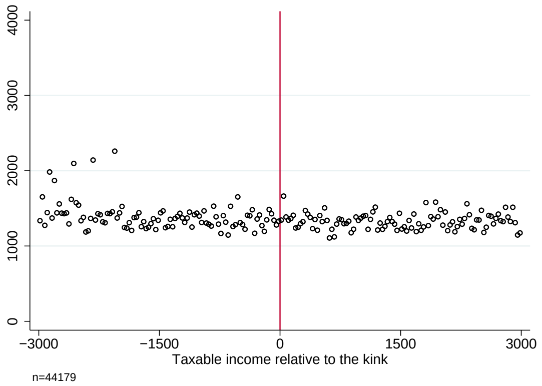

(a)

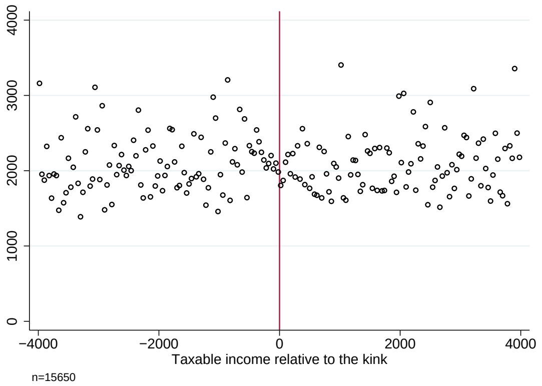

(b)

Primary earners

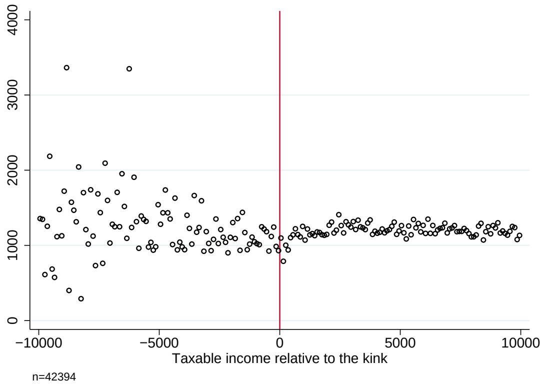

(c)

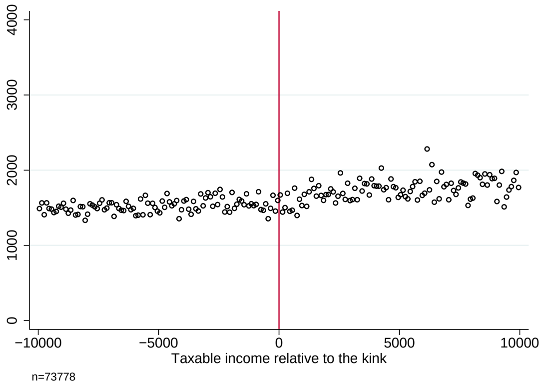

(d)

Secondary earners

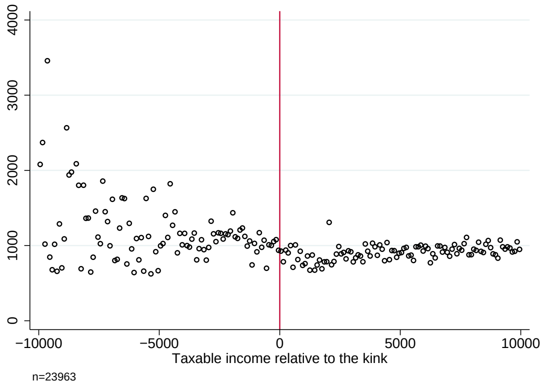

(e)

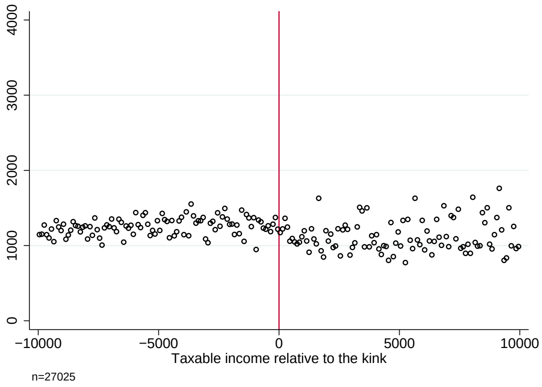

(f)

Notes: Own calculations based on tax return data from Statistics Netherlands.

Figure A9: Using gross income of the primary earner instead of taxable income shows no bunching around the kink

Kink 1

Kink 2

Density around the kink

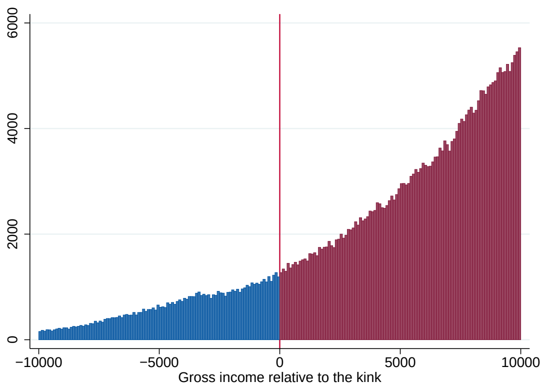

(a)

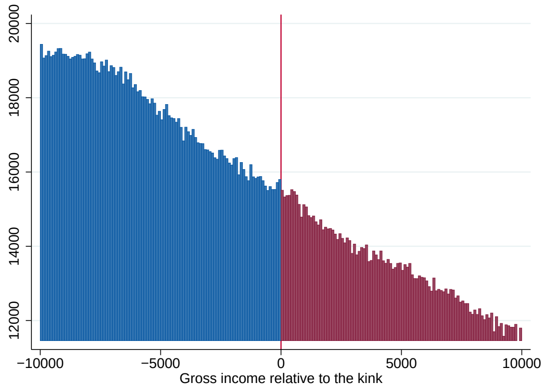

(b)

Notes: Own calculations based on tax return data from Statistics Netherlands.

Figure A10: Characteristics of primary earners with gross income relative to the kink

Kink 1 | Kink 2

Female

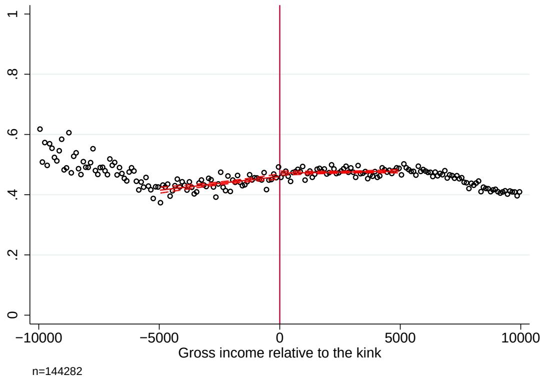

(a)

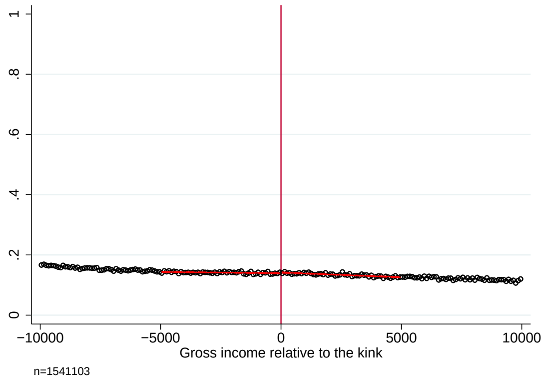

(b)

Age

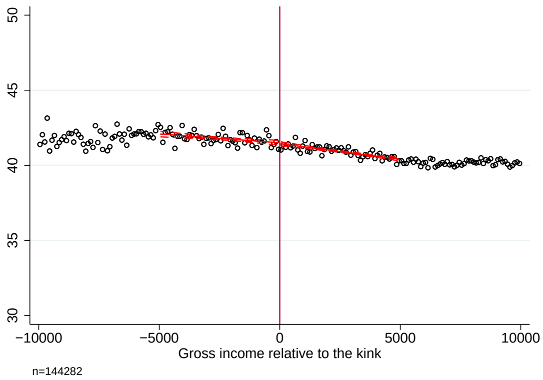

(c)

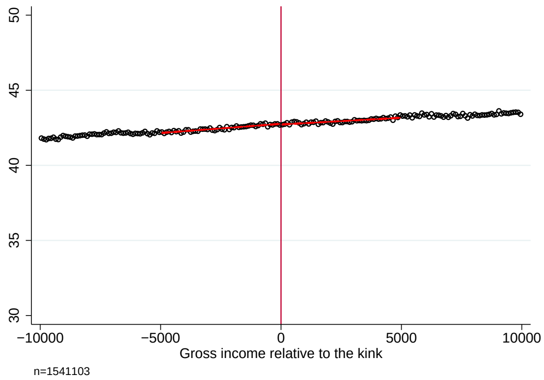

(d)

Notes: Own calculations based on register data from Statistics Netherlands. The regression lines are linear functions without any control variables, with a separate intercept and slope on the right-hand side of the kink.

Figure A10: Characteristics of primary earners with gross income relative to the kink (cont.)

Kink 1 Kink 2

Foreign

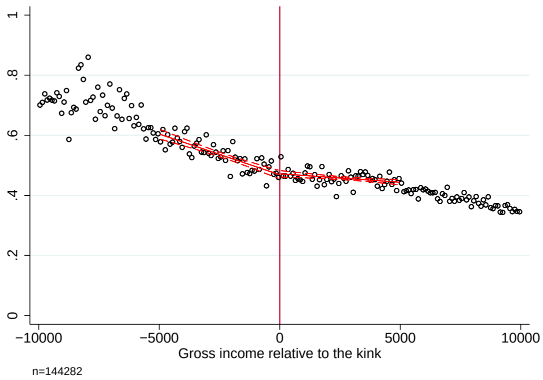

(e)

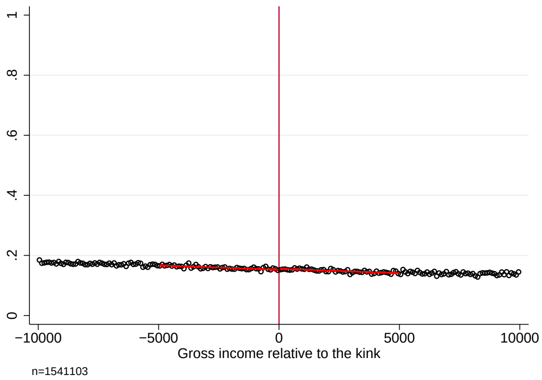

(f)

Number of children

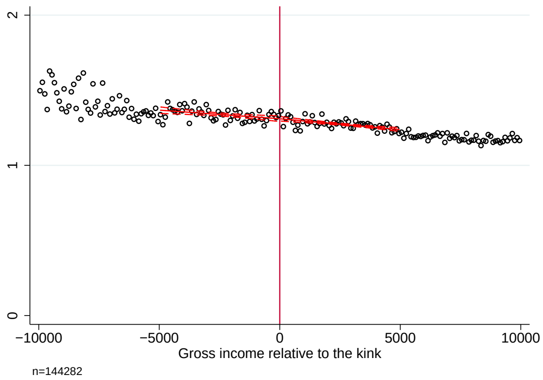

(g)

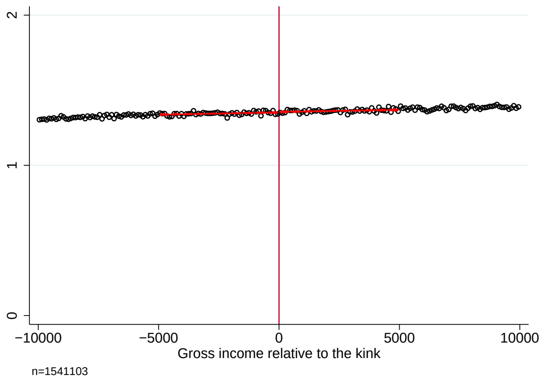

(h)

Notes: Own calculations based on register data from Statistics Netherlands. The regression lines are linear functions without any control variables, with a separate intercept and slope on the right-hand side of the kink.

Figure A11: Total effective marginal tax rates (solid black lines) and decomposition for childless singles and lone parents at kink 1

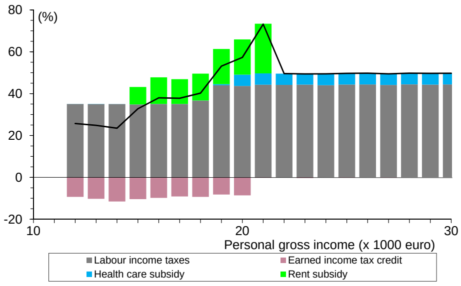

(a) Childless singles

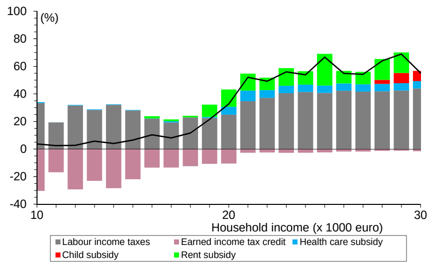

(b) Lone parents

Notes: Own calculations using the tax-benefit microsimulation model MIMOSI of CPB Netherlands Bureau for Economic Policy Analysis (Koot et al., 2016).

Figure A12: Own deducted amount in 2012 and 2013

2012 | 2013

Singles

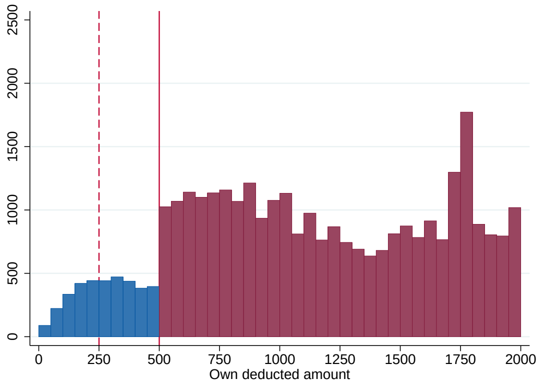

(a)

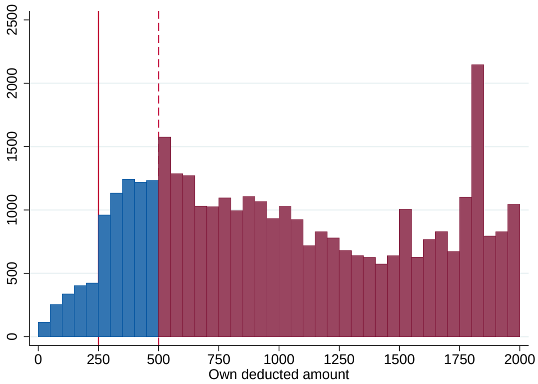

(b)

Primary earners

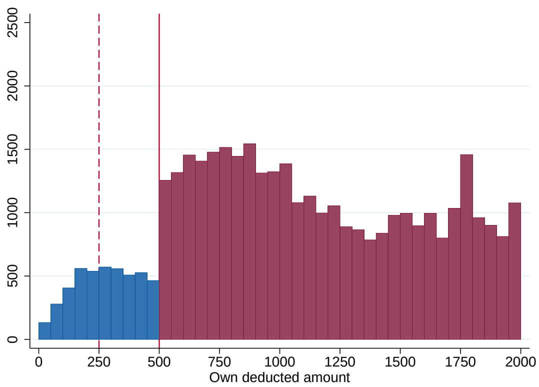

(c)

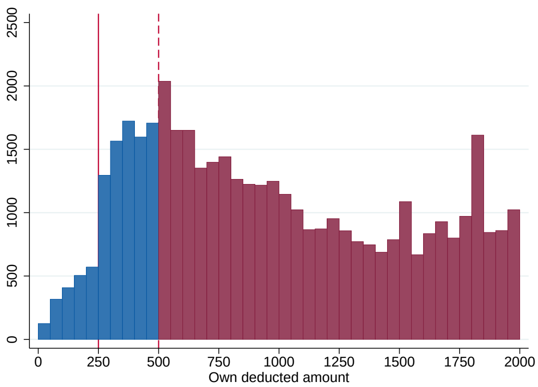

(d)

Secondary earners

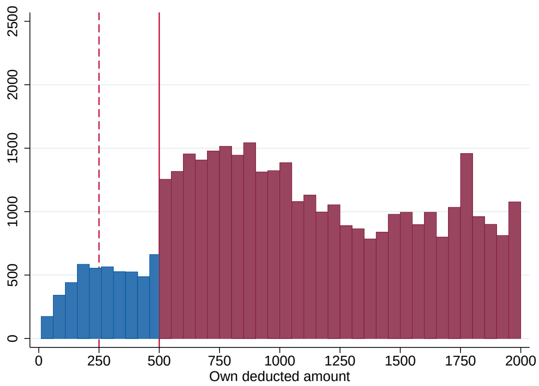

(e)

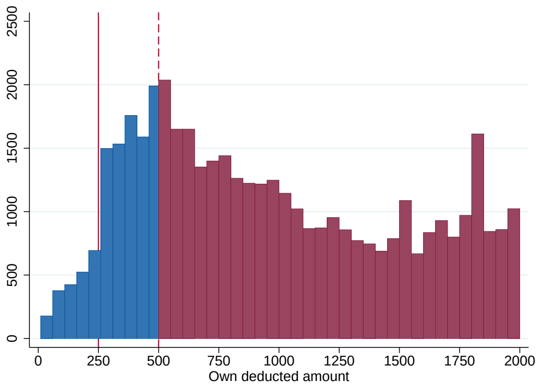

(f)

Notes: Own calculations based on tax return data from Statistics Netherlands.

CHAPTER 6

Summary and Conclusions

Chapter 2. Automatic reaction: what happens to workers at firms that automate?

Chapter 2 considers the impact of automation within the firm on workers' outcomes. We use data on 36,085 Dutch firms employing close to five million workers on average per year in the period 2000–2016. We compare workers at firms that automate with workers at firms that automate later. Our analysis shows that automation at the firm increases the probability to separate from the firm by around 24% for incumbent workers (those who have been at the firm for at least three years). For them, firm separation is followed by an increase in time spent in unemployment. Due to the increased incidence of unemployment, workers experience on average a decline in cumulative wage income of around 8% of yearly earnings after five years. Wage rates however, do not appear to be affected. We also find that earnings losses are larger for older workers and are only partially offset by benefits. While the probability to separate from the firm also increases for those recently hired by the firm (one or two years before the automation event), these workers experience no decline in income.

Chapter 3. Bad start, bad match? The early career effects of graduating in a recession for vocational and academic graduates

Chapter 3 estimates the effect of graduating in a recession for high-educated workers in the Netherlands. Using data on all Dutch graduates between 1996 and 2012, I find that academic graduates suffer initial wage losses of 10% for each percentage point decline (around half of a standard deviation) in field-specific employment at graduation. The wage losses fade out after five years on the labor market. The wage losses for vocational graduates are smaller at 6% for each percentage point decline in field-specific employment at graduation. However, they remain persistent at 1% up to at least 8 years after graduation. Employment probabilities for both groups are negatively affected in the first four years on the labor market. Job mobility plays a critical role in recovering from initial wage losses. Both groups are more likely to switch firms and sectors, and when they do switch, they gain more than their counterparts who started in a boom. Switching also resolves sectoral mismatch for academic, but not for vocational graduates.

Chapter 4. Do parents work more when children start school? Evidence from the Netherlands

Chapter 4 examines how parents adjust their working hours when their youngest child goes to school. We show theoretically that there are two effects. First, parents who used to take care of their children during school hours have an increase in time available and are hence expected to increase their working hours. Second, parents whose children attended paid childcare before going to school might decrease their working hours, because they spend less on childcare expenses. Using data on all Dutch parents between 2006 and 2016, we show significant differences in the responses between fathers and mothers. Dutch mothers on average experience an increase in their free time of 13 hours a week when their youngest child goes to school, yet the average number of hours worked per week increases by 0.5 hours after two years. Dutch fathers, who usually already worked full-time, also show a small increase in hours worked of about 0.3 hours.

Chapter 5. Using tax deductions to promote lifelong learning: real and shifting responses

Finally, chapter 5 considers whether a tax deduction stimulates investment in lifelong learning. Workers are allowed to deduct expenses for lifelong learning at their marginal tax rate. We exploit two jumps in the marginal tax rate in a regression kink design to estimate the effect of the deduction. These jumps create exogenous variation in the effective costs of lifelong learning for people with very similar income levels. We apply this method to data on the universe of Dutch tax payers between 2006 and 2012. For singles we find heterogeneous effects. For low-income singles we find no effect of the lower costs of lifelong learning due to the jump in the marginal tax rate. However, for high-income singles we find a 10% increase in the probability to file lifelong learning expenditures. We find that these effects are primarily driven by higher-educated middle-aged males. For couples we find small effects for primary earners and no effects for secondary earners.

Samenvatting (Dutch Summary)

Schokken op de arbeidsmarkt kunnen van grote invloed zijn op het leven en de carrières van mensen. Bij een schok kun je bijvoorbeeld denken aan een fabriek die moet sluiten in het midden van een recessie. Mensen verliezen hierdoor, buiten hun eigen schuld om. Desalniettemin kunnen de gevolgen groot zijn. Hoewel veel mensen relatief snel weer nieuw werk vinden, kan het voor sommigen het begin zijn van een lange periode van werkloosheid, een langdurig lager inkomen (Jacobson et al., 1993) en slechtere gezondheid (Rege et al., 2009). Onderzoek laat zelfs zien dat werknemers die te maken hebben met onvrijwillig baanverlies gemiddeld eerder overlijden (Sullivan and von Wachter, 2009). Ook doen hun kinderen het slechter in school (Rege et al., 2011).

Schokken op de arbeidsmarkt kunnen veel verschillende oorzaken hebben. Een belangrijke oorzaak is de conjunctuur, zoals bij het hierboven genoemde voorbeeld. Maar ook instuties en veranderingen in beleid kunnen leiden tot schokken. Structurele veranderingen, zoals globalisering en de introductie van nieuwe technologie in het productieproces kunnen ook van grote invloed zijn op het werk van mensen. De eerste drie hoofdstukken van dit proefschrift onderzoeken hoe drie verschillende type schokken invloed hebben op de arbeidsparticipatie, het inkomen en het gebruik van sociale zekerheid van mensen. De hoofdstukken gaan respectievelijk over schokken veroorzaakt door automatisering, de conjunctuur en door instituties. Er zijn verschillende beleidsreacties mogelijk op dit soort schokken. Vaak bestaat de reactie uit een compensatie voor de werknemers die geraakt worden, zoals via een werkloosheidsuitkering. Dit is echter slechts een tijdelijk antwoord om het directe inkomensverlies op te vangen. Mensen die hun baan verliezen doordat bijvoorbeeld een robot hun werk heeft overgenomen, hebben vaak nieuwe vaardigheden nodig om nieuw werk te vinden. Een veel gehoorde maatregel is dan om te investeren in training. Het is echter nog niet duidelijk hoe dit soort beleid het beste vorm gegeven kan worden. In hoofdstuk 5 onderzoeken we of een fiscale aftrekpost effectief is bij het stimuleren van training.

In hoofdstuk 2 onderzoeken we de impact van automatisering binnen een bedrijf op de uitkomsten van werknemers bij dat bedrijf. We gebruiken data over 36.085 Nederlandse bedrijven met bijna vijf miljoen werknemers per jaar voor de periode 2000–2016. We vergelijken werknemers bij bedrijven die automatiseren met werknemers bij bedrijven die later automatiseren. Onze analyse laat zien dat automatisering binnen een bedrijf ertoe leidt dat werknemers die al minimaal drie jaar bij het bedrijf werken voordat het gaat automatiseren een hogere kans hebben om te vertrekken bij het bedrijf na automatisering. Na vertrek worden ze vaker werkloos. Ze verliezen cumulatief ongeveer 11% van een jaarinkomen na vijf jaar door automatisering. We vinden geen effecten op het loon voor deze werknemers als ze werk hebben. Het inkomensverlies komt dus puur doordat ze vaker werkloos zijn dan werknemers die niet met automatisering te maken hebben. De sociale zekerheid compenseert maar een deel van het inkomensverlies. We laten ook zien dat grote investeringen in computers niet tot dit soort negatieve gevolgen leiden voor werknemers.

Hoofdstuk 3 onderzoekt de gevolgen van afstuderen tijdens een recessie voor mensen die een hbo- of wo-opleiding hebben afgerond. Ik gebruik data over alle afgestudeerden tussen 1996 en 2012. Zowel hbo- als wo-opgeleiden die in een recessie starten verdienen initieel minder dan hun leeftijdgenoten die op een beter moment starten. Voor wo-opgeleiden is het initiële loon 10% lager per procentpunt afname in de werkgelegenheid in hun studierichting (ongeveer een halve standaardafwijking). Na vijf jaar halen ze het verlies in. Voor hbo opgeleiden is het initiële verlies met 6% kleiner, maar ook na 8 jaar verdienen zij nog 1% minder per procentpunt afname in de werkgelegenheid in hun studierichting dan mensen die in een hoogconjunctuur zijn gestart. De kans op werk is voor beide groepen lager in de eerste vier jaar na afstuderen. Ze halen het verlies in door vaker te wisselen van baan en sector. En als ze wisselen, maken ze ook grotere loonstappen dan hun leeftijdgenoten die op een beter moment zijn gestart. Hbo afgestudeerden blijven echter ook na 8 jaar nog vaker in een mismatch.

In hoofdstuk 4 onderzoeken we hoe ouders hun gewerkte uren aanpassen als hun jongste kind naar school gaan. Theoretisch zijn er twee effecten te verwachten van deze verandering. Ten eerste, voor ouders die zelf voor hun kind zorgden tijdens schooluren komt er meer tijd beschikbaar. Dit kan leiden tot een toename in hun gewerkte uren. Ten tweede, ouders die gebruik maakten van betaalde kinderopvang hoeven daar nu minder aan uit te geven, en gaan daardoor mogelijk minder werken. We gebruiken gegevens over alle ouders tussen 2006 en 2016. Moeders gaan gemiddeld

0,5 uur per week extra werken als hun jongste kind naar school gaat, hoewel zij tegelijkertijd gemiddeld 13 uur per week minder tijd kwijt zijn aan het zorgen voor hun kind. Vaders laten een kleinere toename zien van ongeveer 0,3 uur per week. Zij werkten echter ook al vaker voltijds voor hun kind naar school gaat.

Tenslotte bekijken we in hoofdstuk 5 of een fiscale aftrekpost ertoe leidt dat werkenden meer investeren in scholing. Werkenden, zowel zzp'ers als werknemers, mogen hun kosten voor scholing aftrekken tegen hun marginale belastingtarief. Het marginale belastingtarief kent twee discrete sprongen. Hierdoor verschillen de effectieve kosten voor scholing voor mensen die net onder en net boven deze grenzen verdienen. We gebruiken een zogenaamd *regression kink design* om deze mensen te vergelijken. Voor onze analyse gebruiken we data over alle belastingbetalers tussen 2006 en 2012. Voor alleenstanden vinden we heterogene effecten. Alleenstaanden met een laag inkomen reageren niet op de lagere kosten voor scholing. Alleenstaanden met een hoog inkomen laten echter een 10% toename zien in de kans dat ze scholingsuitgaven aftrekken. Deze extra uitgaven worden vooral gedaan door hoger opgeleide mannen van middelbare leeftijd. Binnen stellen vinden we kleine positieve effecten voor de meestverdiener en geen effect voor de minstverdiener.

About the Author

Adrianus Willem (Wiljan) van den Berge (1987) obtained a bachelor's degree in Economics and a bachelor's degree in Philosophy of Science at Erasmus University Rotterdam in 2009. He then started the research master in Philosophy of Economics and obtained his degree in 2012 at Erasmus University Rotterdam. Finally, he obtained a master's degree in Economics at Tilburg University in 2013. Wiljan started as an intern at the CPB Netherlands Bureau for Economic Policy Analysis (*Centraal Planbureau*) in the Spring of 2013 and wrote his master's thesis at the CPB. After graduation he started working at the CPB in the department of Labour, Education and Pensions. There he contributed to many research and policy projects in labour and education, covering topics such as the impact of new technology on the labour market, the regional diffusion of knowledge and displacement on the labour market. He started as an external PhD student first at Tilburg University in 2014 supervised by prof. dr. Jan van Ours and then moved to Erasmus University Rotterdam. In 2018 and 2019 he spent time as a visting researcher at Boston University.

Portfolio

Courses during PhD

Machine Learning and Data science	2018
Applied Microeconometrics II (Tinbergen Institute)	2015
Applied Microeconometrics I (Tinbergen Institute)	2014

Presentations at conferences and seminars

European Association of Labour Economists conference	2019
IZA Summer School	2019
Erasmus University Rotterdam Brown Bag seminar	2019
Groningen University seminar	2019
Dutch Economists Day conference	2018
Boston University TPRI lunch seminar	2018
Bentley College seminar	2018
Utrecht University Day of Digital Times seminar	2018
European Society for Population Economics conference	2017
European Association of Labour Economists conference	2017
Maastricht University seminar	2017
IZA Economics of Education workshop	2016
Tilburg University seminar	2016
Dutch Economists Day conference	2013

Teaching and supervision

Guest lecture and assignment on “Technology and the labour market” in MSc course Labour Economics, Tilburg University	2017 - 2019
Supervision of MSc thesis Emiel van Bezooijen at CPB	2019
Supervision of MSc thesis Lisa van Bentum at CPB	2018
Supervision of MSc student intern Johanna Lenitz at CPB	2017
Supervision of MSc student interns Charlotte Crooijmans, Nina Pkhikidze and Ronja Rottger at CPB	2015

Research visits

Boston University TPRI	2019
Boston University TPRI	2018

Bibliography

Abeler, J. and Marklein, F. (2016). Fungibility, Labels, and Consumption. *Journal of the European Economic Association*, 15(1):99–127.

Abramovsky, L., Battistin, E., Fitzsimons, E., Goodman, A., and Simpson, H. (2011). Providing Employers with Incentives to Train Low-Skilled Workers: Evidence from the UK Employer Training Pilots. *Journal of Labor Economics*, 29(1):153–193.

Acemoglu, D. and Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. *Handbook of Labor Economics*, 4:1043–1171.

Acemoglu, D. and Pischke, J.-S. (1998). Why Do Firms Train? Theory and Evidence. *The Quarterly Journal of Economics*, 113(1):79–119.

Acemoglu, D. and Restrepo, P. (2018a). Artificial Intelligence, Automation and Work. In Agrawal, A. K., Gans, J., and Goldfarb, A., editors, *The Economics of Artificial Intelligence*. University of Chicago Press.

Acemoglu, D. and Restrepo, P. (2018b). Modeling Automation. *AEA Papers and Proceedings*, 108:48–53.

Acemoglu, D. and Restrepo, P. (2018c). Robots and Jobs: Evidence from US Labor Markets. MIT.

Acemoglu, D. and Restrepo, P. (2018d). The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment. *American Economic Review*, 108(6):1488–1542.

Altonji, J. G., Kahn, L. B., and Speer, J. D. (2016). Cashier or Consultant? Entry Labor Market Conditions, Field of Study and Career Success. *Journal of Labor Economics*, 34(S1):361–401.

Angrist, J. D. and Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Autor, D. and Salomons, A. (2018). Is Automation Labor-Displacing? Productivity Growth, Employment, and the Labor Share. *Brookings Papers on Economic Activity*, Spring.

Autor, D. H. (2015). Why Are There Still So Many Jobs? The History and Future of Workplace Automation. *Journal of Economic Perspectives*, 29(3):3–30.

Autor, D. H., Dorn, D., Hanson, G. H., and Song, J. (2014). Trade Adjustment: Worker-Level Evidence. *The Quarterly Journal of Economics*, 129(4):1799–1860.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The Skill Content Of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333.

Azoulay, P., Graff Zivin, J. S., and Wang, J. (2010). Superstar Extinction. *The Quarterly Journal of Economics*, 125(2):549–589.

Barreca, A. I., Guldi, M., Lindo, J. M., and Waddell, G. R. (2011). Saving Babies? Revisiting the Effect of Very Low Birth Weight Classification. *The Quarterly Journal of Economics*, 126(4):2117–2123.

Barreca, A. I., Lindo, J. M., and Waddell, G. R. (2016). Heaping-induced Bias in Regression-discontinuity Designs. *Economic Inquiry*, 54(1):268–293.

Beaudry, P. and DiNardo, J. (1991). The Effect of Implicit Contracts on the Movement of Wages over the Business Cycle. *Journal of Political Economy*, 99(4):665–688.

Beiler, H. (2017). Do You Dare? The Effect of Economic Conditions on Entrepreneurship Among College Graduates. *Labour Economics*, 47:64 – 74.

Benzell, S. G., Kotlikoff, L. J., LaGarda, G., and Sachs, J. D. (2016). Robots Are Us: Some Economics of Human Replacement. NBER Working Paper 20941, National Bureau of Economic Research, Inc.

Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How Much Should We Trust Differences-In-Differences Estimates? *The Quarterly Journal of Economics*, 119(1):249–275.

Bessen, J. (1999). Real Options and the Adoption of New Technologies. Online working paper at http://www. researchoninnovation. org/.

Bessen, J., Goos, M., Salomons, A., and van den Berge, W. (2019). Automatic Reaction: What Happens to Workers at Firms that Automate? *CPB Discussion Paper No. 390.*

Bessen, J., Goos, M., Salomons, A., and van den Berge, W. (2020). Evidence on Firms' Automation Activities. *American Economic Association Papers & Proceedings*, Forthcoming.

Blackwell, M., Iacus, S., King, G., and Porro, G. (2009). CEM: Coarsened Exact Matching in Stata. *Stata Journal*, 9(4):524–546.

Blau, F. D. and Kahn, L. M. (2007). Changes in the Labor Supply Behavior of Married Women: 1980–2000. *Journal of Labor Economics*, 25(3):393–438.

Blundell, R., Bozio, A., and Laroque, G. (2013). Extensive and Intensive Margins of Labour Supply: Work and Working Hours in the US, the UK and France. *Fiscal Studies*, 34(1):1–29.

Boskin, M. J. (1975). Notes on the Tax Treatment of Human Capital. Working Paper 116, National Bureau of Economic Research, Cambridge, MA.

Brier, G. W. (1950). Verification of Forecasts Expressed in Terms of Probability. *Monthly Weather Review*, 78:1–3.

Brunner, B. and Kuhn, A. (2014). The Impact of Labor Market Entry Conditions on Initial Job Assignment and Wages. *Journal of Population Economics*, 27(3):705–738.

Bulman, G. B. and Hoxby, C. M. (2015). The Returns to the Federal Tax Credits for Higher Education. In *Tax Policy and the Economy, Volume 29*, pages 13–88. University of Chicago Press, Chicago, IL.

Card, D., Johnston, A., Leung, P., Mas, A., and Pei, Z. (2015a). The Effect of Unemployment Benefits on the Duration of Unemployment Insurance Receipt: New Evidence from a Regression Kink Design in Missouri, 2003–2013. *The American Economic Review*, 105(5):126–130.

Card, D., Lee, D. S., Pei, Z., and Weber, A. (2015b). Inference on Causal Effects in a Generalized Regression Kink Design. *Econometrica*, 83(6):2453–2483.

Carneiro, A., Portugal, P., and Raposo, P. (2015). Decomposing the Wage Losses of Displaced Workers: The Role of the Reallocation of Workers into Firms and Job Titles. IZA Discussion Papers 9220, Institute for the Study of Labor (IZA).

Cascio, E. U. (2009). Maternal Labor Supply and the Introduction of Kindergartens into American Public Schools. *Journal of Human Resources*, 44(1):140–170.

Cattaneo, M. D., Jansson, M., and Ma, X. (2016). Simple Local Regression Distribution Estimators with an Application to Manipulation Testing. University of Michigan.

Cattaneo, M. D., Jansson, M., and Ma, X. (2017). rddensity: Manipulation Testing based on Density Discontinuity. University of Michigan.

Chetty, R. (2009). Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance. *American Economic Journal: Economic Policy*, 1(2):31–52.

Chetty, R., Friedman, J. N., and Saez, E. (2013). Using Differences in Knowledge Across Neighborhoods to Uncover the Impacts of the EITC on Earnings. *The American Economic Review*, 103(7):2683–2721.

Cockx, B. and Ghirelli, C. (2016). Scars of Recessions in a Rigid Labor Market. *Labour Economics*, 41:162 – 176.

Cortés, G. (2016). Where Have the Middle-Wage Workers Gone? A Study of Polarization using Panel Data". *Journal of Labor Economics*, forthcoming.

Couch, K. A. and Placzek, D. W. (2010). Earnings Losses of Displaced Workers Revisited. *American Economic Review*, 100(1):572–89.

Dauth, W., Findeisen, S., Südekum, J., and Wößner, N. (2017). German Robots – The Impact of Industrial Robots on Workers. Technical report, Institute for Employment Research, Nuremberg, Germany.

Dauth, W., Findeisen, S., Südekum, J., and Wößner, N. (2018). Adjusting to Robots: Worker Level Evidence . Working paper.

Davis, S. J. and Von Wachter, T. (2011). Recessions and the Costs of Job Loss. *Brookings Papers on Economic Activity*, 42(2):1–71.

Deelen, A., de Graaf-Zijl, M., and van den Berge, W. (2018). Labour Market Effects of Job Displacement for Prime-age and Older Workers. *IZA Journal of Labor Economics*, 7(1):3.

Del Bono, E. and Vuri, D. (2011). Job Mobility and the Gender Wage Gap in Italy. *Labour Economics*, 18(1):130 – 142.

Devereux, P. J. and Hart, R. A. (2010). Forced to be Rich? Returns to Compulsory Schooling in Britain. *The Economic Journal*, 120(549):1345–1364.

Dinlersoz, E. and Wolf, Z. (2018). Automation, Labor Share, and Productivity: Plant-Level Evidence from U.S. Manufacturing. Working Papers 18-39, Center for Economic Studies, U.S. Census Bureau.

Dixon, J., Hong, B., and Wu, L. (2019). The Employment Consequences of Robots: Firm-Level Evidence. Ssrn discussion paper, SSRN.

Doerrenberg, P., Peichl, A., and Siegloch, S. (2017). The Elasticity of Taxable Income in the Presence of Deduction Possibilities. *Journal of Public Economics*, 151:41–55.

Doms, M., Dunne, T., and Troske, K. R. (1997). Workers, Wages, and Technology. *The Quarterly Journal of Economics*, 112(1):253–290.

Doms, M. E. and Dunne, T. (1998). Capital Adjustment Patterns in Manufacturing Plants. *Review of Economic Dynamics*, 1(2):409–429.

Donald, S. G. and Lang, K. (2007). Inference with Difference-in-Differences and Other Panel Data. *The Review of Economics and Statistics*, 89(2):221–233.

Duggan, M., Garthwaite, C., and Goyal, A. (2016). The Market Impacts of Pharmaceutical Product Patents in Developing Countries: Evidence from India. *American Economic Review*, 106(1):99–135.

Dustmann, C. and Schönberg, U. (2012). What Makes Firm-Based Vocational Training Schemes Successful? The Role of Commitment. *American Economic Journal: Applied Economics*, 4(2):36–61.

Dynarski, S. and Scott-Clayton, J. (2016). Tax Benefits for College Attendance. Working Paper 22127, National Bureau of Economic Research, Cambridge, MA.

Eaton, J. and Rosen, H. S. (1980). Taxation, Human Capital, and Uncertainty. *The American Economic Review*, 70(4):705–715.

Edin, P.-A., Evans, T., Graetz, G., Hernnäs, S., and Michaels, G. (2019). Individual Consequences of Occupational Decline. Working paper.

Erpelinck, K. and Van Sonsbeek, J.-M. (2012). Economische criss achtervolgt jonge hoogopgeleiden. *Economische en Statistische Berichten*, 97(4632):198–201.

Eurobarometer (2017). Attitudes Towards the Impact of Digitisation and Automation on Daily Life. Technical report, European Commission, Directorate-General for Communication, Special Eurobarometer 460.

Eurostat (2012). Being Young in Europe Today: Education.

Eurostat (2016). Adult Education Survey.

Fadlon, I. and Nielsen, T. H. (2017). Family Health Behaviors. Working Paper 24042, National Bureau of Economic Research.

Fernández-Kranz, D. and Rodríguez-Planas, N. (2018). The Perfect Storm: Graduating during a Recession in a Segmented Labor Market. *ILR Review*, 71(2):492–524.

Fersterer, J., Pischke, J.-S., and Winter-Ebmer, R. (2008). Returns to Apprenticeship Training in Austria: Evidence from Failed Firms. *The Scandinavian Journal of Economics*, 110(4):733–753.

Finseraas, H., Hardoy, I., and Schøne, P. (2016). School Enrolment and Mothers' Labor Supply: Evidence from a Regression Discontinuity Approach. *Review of Economics of the Household*, pages 1–18.

Fisher, S. R. A. (1935). *The Design of Experiments*. Macmillan.

Fitzenberger, B., Licklederer, S., and Zwiener, H. (2015). Mobility Across Firms and Occupations Among Graduates from Apprenticeship. *Labour Economics*, 34:138 – 151.

Fitzpatrick, M. D. (2010). Preschoolers Enrolled and Mothers at Work? The Effects of Universal Prekindergarten. *Journal of Labor Economics*, 28(1):51–85.

Fitzpatrick, M. D. (2012). Revising Our Thinking About the Relationship Between Maternal Labor Supply and Preschool. *Journal of Human Resources*, 47(3):583–612.

Fouarge, D. (2009). Effecten van crisis voor schoolverlaters. *ESB*, 94(4568):554–556.

Ganong, P. and Jäger, S. (2018). A Permutation Test for the Regression Kink Design. *Journal of the American Statistical Association, forthcoming.*

Gelbach, J. B. (2002). Public Schooling for Young Children and Maternal Labor Supply. *The American Economic Review*, 92(1):307–322.

Genda, Y., Kondo, A., and Ohta, S. (2010). Long-Term Effects of a Recession at Labor Market Entry in Japan and the United States. *Journal of Human Resources*, 45(1):157–196.

Gibbons, R. and Waldman, M. (2004). Task-Specific Human Capital. *The American Economic Review*, 94(2):203–207.

Gibbons, R. and Waldman, M. (2006). Enriching a Theory of Wage and Promotion Dynamics inside Firms. *Journal of Labor Economics*, 24(1):59–107.

Göggel, K. and Zwick, T. (2012). Heterogeneous Wage Effects of Apprenticeship Training. *The Scandinavian Journal of Economics*, 114(3):756–779.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *American Economic Review*, 104(8):2509–2526.

Görlitz, K. (2010). The Effect of Subsidizing Continuous Training Investments: Evidence from German Establishment Data. *Labour Economics*, 17(5):789 – 798.

Görlitz, K. and Tamm, M. (2016). The Returns to Voucher-Financed Training on Wages, Employment and Job Tasks. *Economics of Education Review*, 52:51 – 62.

Goux, D. and Maurin, E. (2010). Public School Availability for Two-Year Olds and Mothers' Labour Supply. *Labour Economics*, 17(6):951 – 962.

Graetz, G. and Michaels, G. (2018). Robots at Work. *Review of Economics and Statistics*, forthcoming.

Grenet, J. (2013). Is Extending Compulsory Schooling Alone Enough to Raise Earnings? Evidence from French and British Compulsory Schooling Laws. *The Scandinavian Journal of Economics*, 115(1):176–210.

Hall, C. (2016). Does More General Education Reduce the Risk of Future Unemployment? Evidence from an Expansion of Vocational Upper Secondary Education. *Economics of Education Review*, 52:251 – 271.

Haltiwanger, J., Cooper, R., and Power, L. (1999). Machine Replacement and the Business Cycle: Lumps and Bumps. *American Economic Review*, 89(4):921–946.

Haltiwanger, J. C., Hyatt, H. R., Kahn, L. B., and McEntarfer, E. (2017). Cyclical Job Ladders by Firm Size and Firm Wage. *NBER Working Paper 23458*.

Hanushek, E. A., Schwerdt, G., Woessmann, L., and Zhang, L. (2017). General Education, Vocational Education, and Labor-Market Outcomes over the Lifecycle. *Journal of Human Resources*, 52(1):48–87.

Hardoy, I. and Schøne, P. (2014). Displacement and Household Adaptation: Insured by the Spouse or the State? *Journal of Population Economics*, 27(3):683–703.

Harris, M. and Holmstrom, B. (1982). A Theory of Wage Dynamics. *The Review of Economic Studies*, 49(3):315–333.

Heckman, J. J., Lochner, L. J., and Todd, P. E. (2006). Earnings Functions, Rates of Return and Treatment Effects: The Mincer Equation and Beyond. volume 1 of *Handbook of the Economics of Education*, pages 307 – 458. Elsevier.

Heijke, H., Meng, C., and Ris, C. (2003). Fitting to the Job: The Role of Generic and Vocational Competencies in Adjustment and Performance. *Labour Economics*, 10(2):215 – 229.

Hershbein, B. (2012). Graduating High School in a Recession: Work, Education, and Home Production. *The B.E. Journal of Economic Analysis & Policy*, 12(1).

Hidalgo, D., Oosterbeek, H., and Webbink, D. (2014). The Impact of Training Vouchers on Low-skilled Workers. *Labour Economics*, 31:117 – 128.

Hoffman, M. and Burks, S. V. (2013). Training Contracts, Worker Overconfidence, and the Provision of Firm-Sponsored General Training. *Working Paper*.

Hoxby, C. M. and Bulman, G. B. (2016). The Effects of the Tax Deduction for Postsecondary Tuition: Implications for Structuring Tax-based Aid. *Economics of Education Review*, 51:23 – 60.

Humburg, M., de Grip, A., and van der Velden, R. (2017). Which Skills Protect Graduates Against a Slack Labour Market? *International Labour Review*, 156(1):25–43.

Iacus, S. M., King, G., and Porro, G. (2012a). Causal Inference Without Balance Checking: Coarsened Exact Matching. *Political Analysis*, 20(1):1–24.

Iacus, S. M., King, G., and Porro, G. (2012b). Causal Inference Without Balance Checking: Coarsened Exact Matching. *Political Analysis*, 20(1):1–24.

Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993). Earnings Losses of Displaced Workers. *The American Economic Review*, pages 685–709.

Kahn, L. B. (2010). The Long-term Labor Market Consequences of Graduating from College in a Bad Economy. *Labour Economics*, 17(2):303 – 316.

Kennedy, P. E. (1995). Randomization Tests in Econometrics. *Journal of Business & Economic Statistics*, 13(1):85–94.

Kinderopvang, B. (2016). Factsheet Kinderopvang.

Kleven, H. J. (2016). Bunching. *Annual Review of Economics*, 8:435–464.

Koch, M., Manuylov, I., and Smolka, M. (2019). Robots and Firms. Working paper.

Kondo, A. (2015). Differential Effects of Graduating During a Recession Across Gender and Race. *IZA Journal of Labor Economics*, 4(1):23.

Koning, P. and Lindeboom, M. (2015). The Rise and Fall of Disability Insurance Enrollment in the Netherlands. *Journal of Economic Perspectives*, 29(2):151–72.

Koot, P., Vlekke, M., Berkhout, E., and Euwals, R. (2016). MIMOSI: Microsimulatiemodel voor belastingen, sociale zekerheid, loonkosten en koopkracht. CPB Background Document, The Hague.

Kopczuk, W. (2005). Tax Bases, Tax Rates and the Elasticity of Reported Income. *Journal of Public Economics*, 89(11–12):2093–2119.

Korpi, T. and Mertens, A. (2003). Training Systems and Labor Mobility: A Comparison between Germany and Sweden. *The Scandinavian Journal of Economics*, 105(4):597–617.

Ladner, P., Looney, A., and Kroft, K. (2009). Salience and Taxation: Theory and Evidence. *The American Economic Review*, 99(4):1145–1177.

Lafortune, J., Rothstein, J., and Schanzenbach, D. W. (2018). School Finance Reform and the Distribution of Student Achievement. *American Economic Journal: Applied Economics*, 10(2):1–26.

Landais, C. (2015). Assessing the Welfare Effects of Unemployment Benefits Using the Regression Kink Design. *American Economic Journal: Economic Policy*, 7(4):243–278.

Lefranc, A. (2003). Labor Market Dynamics and Wage Losses of Displaced Workers in France and the United States. *SSRN Electronic Journal*.

Leuven, E. and Oosterbeek, H. (2004). Evaluating the Effect of Tax Deductions on Training. *Journal of Labor Economics*, 22(2):461–488.

Leuven, E. and Oosterbeek, H. (2012). The Responsiveness of Training Participation to Tax Deductibility. *Working Paper*.

Leuven, E., Oosterbeek, H., Sloof, R., and Van Klaveren, C. (2005). Worker Reciprocity and Employer Investment in Training. *Economica*, 72(285):137–149.

Liu, K., Salvanes, K. G., and Sï¿½rensen, E. . (2016). Good Skills in Bad Times: Cyclical Skill Mismatch and the Long-term Effects of Graduating in a Recession. *European Economic Review*, 84:3–17.

Machin, S., Marie, O., and Vujić, S. (2011). The Crime Reducing Effect of Education. *The Economic Journal*, 121(552):463–484.

Malamud, O. and Pop-Eleches, C. (2010). General Education versus Vocational Training: Evidence from an Economy in Transition. *The Review of Economics and Statistics*, 92(1):43–60.

Malcomson, J. M. (1997). Contracts, Hold-Up, and Labor Markets. *Journal of Economic Literature*, 35(4):1916–1957.

Malcomson, J. M. (1999). Chapter 35 Individual Employment Contracts. volume 3, Part B of *Handbook of Labor Economics*, pages 2291 – 2372. Elsevier.

Martins, P. S., Solon, G., and Thomas, J. P. (2012). Measuring What Employers Do about Entry Wages over the Business Cycle: A New Approach. *American Economic Journal: Macroeconomics*, 4(4):36–55.

McCrary, J. (2008). Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test. *Journal of Econometrics*, 142(2):698 – 714.

McLaughlin, K. J. and Bils, M. (2001). Interindustry Mobility and the Cyclical Upgrading of Labor. *Journal of Labor Economics*, 19(1):94–135.

Miller, C. (2017). The Persistent Effect of Temporary Affirmative Action. *American Economic Journal: Applied Economics*, 9(3):152–90.

Mooi-Reci, I. and Ganzeboom, H. B. (2015). Unemployment Scarring by Gender: Human Capital Depreciation or Stigmatization? Longitudinal Evidence from the Netherlands, 1980–2000. *Social Science Research*, 52:642–658.

Moscarini, G. and Postel-Vinay, F. (2008). The Timing of Labor Market Expansions: New Facts and a New Hypothesis. *NBER Macroeconomics Annual*, 23(1):1–52.

Moscarini, G. and Postel-Vinay, F. (2012). The Contribution of Large and Small Employers to Job Creation in Times of High and Low Unemployment. *American Economic Review*, 102(6):2509–39.

Moscarini, G. and Postel-Vinay, F. (2013). Stochastic Search Equilibrium. *The Review of Economic Studies*, 80(4):1545–1581.

Moscarini, G. and Postel-Vinay, F. (2016). Did the Job Ladder Fail after the Great Recession? *Journal of Labor Economics*, 34(S1):S55–S93.

Nilsen, O. A. and Schiantarelli, F. (2003). Zeros and Lumps in Investment: Empirical Evidence on Irreversibilities and Nonconvexities. *The Review of Economics and Statistics*, 85(4):1021–1037.

Oosterbeek, H. and Webbink, D. (2007). Wage Effects of an Extra Year of Basic Vocational Education. *Economics of Education Review*, 26(4):408 – 419.

Oreopoulos, P., von Wachter, T., and Heisz, A. (2012). The Short- and Long-Term Career Effects of Graduating in a Recession. *American Economic Journal: Applied Economics*, 4(1):1–29.

Oyer, P. (2006). Initial Labor Market Conditions and Long-Term Outcomes for Economists. *Journal of Economic Perspectives*, 20(3):143–160.

Oyer, P. (2008). The Making of an Investment Banker: Stock Market Shocks, Career Choice, and Lifetime Income. *The Journal of Finance*, 63(6):2601–2628.

Pew (2017). Automation in Everyday Life. Technical report, Pew Research Center.

Pindyck, R. (1991). Irreversibility, Uncertainty, and Investment. *Journal of Economic Literature*, 29(3):1110–48.

Rege, M., Telle, K., and Votruba, M. (2009). The Effect of Plant Downsizing on Disability Pension Utilization. *Journal of the European Economic Association*, 7(4):754–785.

Rege, M., Telle, K., and Votruba, M. (2011). Parental Job Loss and Children's School Performance. *The Review of Economic Studies*, 78(4):1462–1489.

Rijksoverheid (2017). Overzicht aantal uren onderwijstijd.

Rothschild, M. (1971). On the Cost of Adjustment. *The Quarterly Journal of Economics*, 85(4):605–622.

Ryan, P. (2001). The School-to-Work Transition: A Cross-National Perspective. *Journal of Economic Literature*, 39(1):34–92.

Saez, E., Slemrod, J., and Giertz, S. (2012). The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review. *Journal of Economic Literature*, 50:1:3–50.

Schwerdt, G., Messer, D., Woessmann, L., and Wolter, S. C. (2012). The Impact of an Adult Education Voucher Program: Evidence From a Randomized Field Experiment. *Journal of Public Economics*, 96(7â€“8):569 – 583.

Shure, N. (2019). School Hours and Maternal Labor Supply. *Kyklos*, 72(1):118–151.

Speer, J. (2016). Wages, Hours, and the School-to-Work Transition: The Consequences of Leaving School in a Recession for Less-Educated Men. *The B.E. Journal of Economic Analysis & Policy*, 16(1):97–124.

Statistics Netherlands (2011). Inpassing van het Nederlande onderwijs in ISCED2011.

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.

Susskind, D. (2017). A Model of Technological Unemployment. Economics Series Working Papers 819, University of Oxford, Department of Economics.

Swart, L., van den Berge, W., and van der Wiel, K. (2019). Do Parents Work More When Children Start School? Evidence from the Netherlands. *CPB Discussion Paper No. 392*.

Thaler, R. (1985). Mental Accounting and Consumer Choice. *Marketing Science*, 4(3):199–214.

Topel, R. H. and Ward, M. P. (1992). Job Mobility and the Careers of Young Men. *Quarterly Journal of Economics*, 107(2):439–479.

Van den Berge, W. (2018). Bad Start, Bad Match? The Early Career Effects of Graduating in a Recession for Vocational and Academic Graduates. *Labour Economics*, 53:75 – 96.

Van den Berge, W., Jongen, E., and van der Wiel, K. (2017). Using Tax Deductions to Promote Lifelong Learning: Real and Shifting Responses. *CPB Discussion Paper No. 353.*

Van der Klaauw, B. and van Vuuren, A. (2010). Job Search and Academic Achievement. *European Economic Review*, 54(2):294 – 316.

Van der Steeg, M. and van Elk, R. (2015). The Effect of Schooling Vouchers on Higher Education Enrollment and Completion of Teachers: A Regression Discontinuity Analysis. CPB Discussion Paper 305, CPB, Den Haag.

Van Ours, J. C. (2009). Jeugdwerkloosheid in barre tijden. *MeJudice*.

Verhaest, D. and Baert, S. (2015). The Early Labour Market Effects of Generally and Vocationally Oriented Higher Education: Is There a Trade-Off? *IZA Discussion Paper No. 9137*.

Von Wachter, T. and Bender, S. (2006). In the Right Place at the Wrong Time: The Role of Firms and Luck in Young Workers' Careers. *American Economic Review*, 96(5):1679–1705.

Webb, M. (2019). The Impact of Artificial Intelligence on the Labor Market. Working paper, Stanford University.

Wolbers, M. H. (2014). Een verloren generatie van jongeren op de arbeidsmarkt? *Tijdschrift voor Arbeidsvraagstukken*, 30(2):103–119.

Wolter, S. C. and Ryan, P. (2011). Chapter 11 - Apprenticeship. volume 3 of *Handbook of the Economics of Education*, pages 521 – 576. Elsevier.

Young, A. (2018). Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results. *The Quarterly Journal of Economics*, 134(2):557–598.