2016-12
Lisbeth Palmhøj Nielsen
PhD Dissertation

Empirical Essays on Child Achievement, Maternal Employment, Parental Leave, and Geographic Mobility

DEPARTMENT OF ECONOMICS AND BUSINESS ECONOMICS
AARHUS BSS • AARHUS UNIVERSITY • DENMARK

SCHOOL OF BUSINESS AND SOCIAL SCIENCES
AARHUS UNIVERSITY

Empirical Essays on Child Achievement, Maternal Employment, Parental Leave and Geographic Mobility

PhD dissertation

Lisbeth Palmhøj Nielsen

Aarhus BSS, Aarhus University
Department of Economics and Business Economics

2016

Table of Contents

Preface

Three years of uplifting freedom and heavy responsibility have gone by, and I find it hard to believe I am now handing in this dissertation. The journey towards this accomplishment has taken me through all states of emotion and all over the world. Many people have helped me reach my goal, and I thank all of you.

I gratefully thank Marianne Simonsen and Nabanita Datta Gupta for all the time and effort you have put into supervising me. Despite the long distance between us, both of you have always found time to read my texts and discuss different aspects of my articles whenever I needed it. Together you have great complementary knowledge and approaches to economics that I have benefitted tremendously from. Even though some months were tough, I never doubted that you wanted what was best for me. You are an inspiration. Thank you for all your hard work.

My time as a research assistant at SFI has been crucial to this dissertation. A special thanks to Anne-Dorthe Hestbæk for always encouraging me and believing in me. The funding from SFI was an extra pad on the shoulder.

Thank you to Miriam Wüst for being the driving force in bringing together a team of researchers to study parental and public interventions in children. Thanks to your efforts and the efforts of the others in the team, we were able to secure funding for the AU RECEIV center.

To my supervisor at SFI Beatrice Schindler-Rangvid, thank you for helping me when I needed it. Great thanks to all my colleagues at SFI for always offering a helping hand, a chat, a coffee, or a beer. I think the work environment at SFI is very special, and I treasure it. When I first started my work in Aarhus, I did not know a soul, but the researchers and PhD-scholars have been welcoming from day one. Thank you for that, and for helpful comments. A special thanks to Susan Stilling for your helpfulness, whether for administrative purposes or correct notation.

I was lucky enough to collaborate with Sean Nicholson and Rachel Dunifon from Cornell University. The joint work with both of you, and Anne Toft Hansen from SFI, has been exciting. Thank you for showing me the craft of writing articles; I have enjoyed our teamwork tremendously. The collaboration brought me to Cornell University for two longer stays, and Sean has visited SFI twice already. I hope for many more visits to come. You have all been open and welcoming, both on email, Skype and face-to-face. A special thanks to Anne for her positive attitude and great effort. You are very talented.

Another advantage from my time as a research assistant was working with Vibeke Myrup Jensen. I gratefully thank you for the opportunity to work with you on the chapter that is now published and part of this dissertation. You have been a role model, a good colleague, and a great collaborator.

Throughout my studies and all the way through my PhD, Shelly Lundberg has been a professional role model. Her models on the family always appealed to me. I emailed her with a pounding heart and asked if I could visit for six months. Instantly she replied that I was welcome. Those six months at University of California, Santa Barbara, stand out as being perfectly seasoned with sun, beautiful beaches, and wonderfully insightful lectures. Shelly, you are welcoming and kind and give great advice. A special thanks to Peter Kuhn for your warm welcome. And thank you to the PhD-students at the Broom Center for Demography for keeping me company during my stay.

My stay at UCSB was more than sandy beaches and research. The love of my life, Max Drasbeck, was living in Atlanta at the time, 4000 kilometers away from me; however, I still managed to become pregnant. Caring for a baby is quite a constraint on work life, but Max, you are the main reason I am handing in my dissertation. You have always been supportive and kept your cool when I could not. Thank you for always standing by me.

Despite the extra worries that come with motherhood, there is nothing like a baby to keep you in the moment and stop worrying about work. Bogatyr, my baby, thank you for keeping me sane during the past two years we have spent together. I adore you.

Parents become more important than ever when you have to care for a child and write a PhD at the same time. Thank you to my mother-in law, Lone Drasbeck, for your always selfless helpfulness. You are the sole reason we make ends meet. Thank you to my own parents and siblings. Even though you live farther away, you help out as much as you can and never miss an opportunity to show your love for us and Bogatyr. We enjoy our relaxing weekends in Roskilde immensely.

Copenhagen, March 2016

Lisbeth Palmhøj Nielsen

Updated preface

I would like to thank the members of the assessment committee, Tor Eriksson (chair), Jane Greve, and Anders Stenberg for their careful reading and insightful comments. I truly appreciate them. Most of the suggestions have been incorporated into the dissertation, while the remaining will be in the near future.

Copenhagen, June 2016
Lisbeth Palmhøj Nielsen

Summary

Despite the obvious overlap of topics, data sources, and econometric methods, this dissertation does not serve as a single unit, but as four self-contained chapters. The overarching theme of the four chapters is research in investments sought to benefit children, both directly and indirectly. Through the work of economist and Nobel laureate James J. Heckman, the literature on early child investments has boomed. The motto is that the earlier the investment, even in utero, the higher the returns. The types of investments are manifold, from direct investments in the child through teaching or care, to indirect investments through parental support, or even geographical interventions.

The first chapter investigates how a direct investment in school children affects their later education. The second chapter takes a step back and explores how parental investments in more or less maternal working hours affect their children's grade point average in the 9^{th} grade. The third chapter explores a parental investment even further from the child, namely subsidies to local geographical areas, and how commuter tax allowance affects labor market outcomes. Even though programs subsidizing geographical areas have proven mostly unsuccessful, some theory suggest that the children of the parents impacted by these reforms benefit from them. The fourth and last chapter explores how parents decide on their time investments in either paid labor or parental leave, and how fathers in families with high bargaining costs experience less parental leave.

The first chapter titled *Increasing the admission rate to upper secondary school: The case of lower secondary school student career guidance* (joint work with Anders Høst and Vibeke Myrup Jensen) examines the effects of career guidance in lower secondary school on upper secondary school admission. The novelty in our study is both the outcome of interest and the intervention itself. Using Danish register data for entire cohorts, we know their educational trajectory. Most articles focus on the grade point average instead of the actual outcome of interest, i.e. whether students continue their education after lower secondary school. Furthermore, as one of only a few studies, we analyze the causal effect of career guidance. We exploit that a Danish school reform changed career guidance from business-as-usual with no or few ground rules to a well-aligned, structured, coherent, and quality-focused career guidance in public, but not in private schools. In a difference-in-difference framework, we estimate that career guidance improved upper secondary school admission for immigrants by 4.0 to 6.3 percentage points, but not for the native Danes. Disregarding the sunk-cost from changing the system, the intervention was cost-neutral, which means that the increase in upper secondary school was close to free.

The second chapter titled *The Effects of Maternal employment on children's Academic Achievement* (joint work with Sean Nicholson, Rachel Dunifon, and Anne Toft Hansen) studies the impact of mothers' work on their children's grade point average (GPA) in the 9th grade. One innovation in the article is that we compare different causal methods and explain the different biases that each method entails. In our preferred model, we use gender and education specific local unemployment as an instrument for local economic activity, and thus the demand for female labor. The first stage of our IV-strategy shows negative, significant signs and high F-test scores, which support the idea that high unemployment in the locality decreases the work hours of the mothers. In the second stage we find that each additional hour per week of maternal employment during the first three years of a child's life is associated with a 1,6 percent of one standard deviation increase in a child's 9th grade GPA. Consequently, children of mothers who work 30 hours per week experience a 48 percent of a standard deviation larger GPA than children of non-working mothers. This finding is in contrast to the literature that largely shows negative effects of mothers' work on their children's cognitive and non-cognitive outcomes. By changing the country from the U.S., that has little support for working parents, to a country that greatly support working parents through subsidized daycare, paid parental leave and sickness leave, the results favor children of working mothers. For policy makers that is a clear indication that supporting working parents benefit not only their children, but in the long run all citizens.

The third chapter titled *Subsidizing rural life: Commuter subsidies and labor market behavior* investigates the effects of increased commuter tax allowance on wages, unemployment, and distance to work. Even though most countries spend large sums of money on the commuter tax allowances, the effects of the subsidy are unknown. This chapter adds to the literature by not only exploring commuter tax allowance, but also deriving the causal effects on different labor market outcomes. The identification stems from a Danish tax reform that increased commuter tax allowance in selected outlying municipalities, but not the rest of the country. Using exact propensity score matching and difference-in-difference, I estimate the average treatment effect on the treated. The results show that residents in the treated areas increased their distance to work, but also that unemployment in the targeted areas was unaffected by the increased commuter tax allowance. Several programs to help outlying areas have proven more or less ineffective.

The fourth chapter titled *The cost of bargaining: The case of fathers' parental leave* discusses the theoretical intra-household bargaining models. I use the gendered spheres threat point instead of the non-cooperative divorce threat point. I conclude that women and men resort to their gendered spheres when they decide on parental leave division, in which case women take up all leave and men none, rather than divorcing if they cannot find a joint solution. I compare the Danish free-choice system, in which parents can freely decide on the parental leave division, to the quota-system introduced in other Scandinavian countries, which earmark a share of the leave for fathers. In the empirical model, I estimate the correlations between different cost of bargaining measures and whether the father takes up leave. Using a linear probability model, I find that in families with high costs, fathers are less likely to take up leave. In two subsample analyses, I furthermore find that the costs of bargaining are more discouraging in families in which the couple has equal bargaining power or the father has higher bargaining power. For more experienced parents, the costs of bargaining also discourage more than for first-time parents. All in all, the chapter concludes that a free-choice system the system distorts cooperation in couples in which the mother wants all leave or the father none.

Resumé (Danish summary)

På trods af det store overlap af emner, datakilder, og økonometriske metoder, skal denne afhandling ikke læses som et samlet værk, men derimod som fire selvstændige kapitler. Det overordnede tema for de fire kapitler omhandler investeringer i børn, både direkte og indirekte. Gennem økonom og nobelprismodtager James J. Heckmans arbejde med tidlig investeringer i børn, har litteraturen på emnet vokset sig stor. Mottoet er, at jo tidligere investering, selv i livmoderen, des højere afkast senere. Litteraturen beskriver mange typer investeringer, fra direkte investeringer i barnet gennem undervisning og omsorg, til indirekte investeringer gennem forældrenes støtte, og selv geografiske tiltag.

I det første kapitel undersøger vi, hvordan en investering i skolebørn påvirker deres senere uddannelse. I det andet kapitel træder vi et skridt tilbage og undersøger, hvordan forældrenes investeringer i mødres kortere eller længere arbejdstid påvirker deres børns karaktergennemsnit i 9. klasse. I det tredje kapitel udforsker jeg en forældre-investering et skridt endnu længere væk fra barnet, nemlig subsidier til geografisk udvalgte kommuner. Jeg undersøger, hvordan et øget befordringsfradrag påvirker arbejdsmarkedsadfærd for de berørte. I det fjerde og sidste kapitel undersøger jeg, hvordan forældrene beslutter hvordan de fordeler deres tid i enten lønnet arbejde eller barselsorlov, og hvordan fædre i familier med høje forhandlingsomkostninger (de emotionelle omkostninger, som der er ved at diskutere med sin partner) tager mindre orlov.

I det første kapitel med titlen *Increasing the admission rate to upper secondary school: The case of lower secondary school student career guidance* (fælles arbejde med Anders Høst og Vibeke Myrup Jensen) finder vi effekten af karrierevejledning i folkeskolen på optagelse og gennemførsel af ungdomsuddannelse. Det nye i vores undersøgelse er, at vi både undersøger ungdomsuddannelse og studievejledning. Ved hjælp af danske registerdata for hele årgange kender vi de unges uddannelsesmæssige veje eller omveje. De fleste artikler fokuserer på karaktergennemsnit fremfor det mere væsentlige, om de studerende fortsætter deres uddannelse efter folkeskolen. Som en af kun få undersøgelser, estimerer vi den kausale effekt af studievejledning. Vi bruger en dansk skolereform, hvor folkeskoler skifter fra *business-as-usual* studievejledning med ingen eller få retningslinjer, til struktureret, sammenhængende og kvalitet-fokuserede studievejledning i folkeskolen. Den nye vejledning blev kun indført i folkeskoler, men ikke i private skoler. Ved brug af den statistiske metode *difference-in-difference* beregner vi at den nye studievejledning har resulteret i mellem 4,0 til 6,3 procentpoint stigning i optag på ungdomsuddannelser, men vi finder ingen effekt for indfødte danskere. Hvis vi ser bort fra de

engangsudgifter, det kostede at omlægge systemet, var reformen af studievejledningen omkostningsneutral. Det betyder, at den strukturerede studievejledning var en tæt på gratis metode til at få flere unge i ungdomsuddannelse.

I det andet kapitel med titlen *The Effects of Maternal employment on children's Academic Achievement* fælles arbejde med Rachel Dunifon, Anne Toft Hansen, og Sean Nicholson) studerer vi effekten af mødres arbejdstid på deres børns karakterer i 9. klasse. En nyskabelse i artiklen er, at vi sammenligner forskellige kausale metoder og forklarer de forskellige fordele og ulemper ved hver metode. I vores foretrukne model, en IV-model, bruger vi lokal arbejdsløshed, stratificeret på køn og uddannelse, som instrument for lokal økonomisk aktivitet, og dermed efterspørgslen efter kvindelig arbejdskraft. Det første trin i vores 2-trins IV-strategi viser en signifikant negativ sammenhæng mellem lokal arbejdsløshed og mødres arbejdstid og høje F-værdier. Dette understøtter ideen om, at den høje arbejdsløshed i lokalområdet nedsætter arbejdstiden for mødrene. Det andet trin viser, at når mødre øger deres arbejdstid med en ekstra time om ugen i løbet af de første tre år af et barns liv stiger deres børns 9. klasses karakterer med 1,6 procent af en standardafvigelse. Derfor vil børn af mødre, der arbejder 30 timer om ugen have 48 procent af en standardafvigelse højere karakterer end børn af ikke-arbejdende kvinder. Dette fund er i modsætning til litteraturen, som i høj grad viser negative effekter af mødres arbejde på deres børns kognitive og ikke-kognitive evner. Ved at undersøge Danmark, et land der i høj grad støtter arbejdende forældre gennem subsidieret barselsorlov, skole, og løn ved sygdom, fremfor USA, der har mindre økonomisk og strukturel støtte til forældre, ændres resultaterne til at fordel for børn af udearbejdende mødre. For de politiske beslutningstagere, er dette et klart tegn på, at støtte til arbejdende forældre er en god investering i fremtiden.

I det tredje kapitel med titlen *Subsidizing rural life: Commuter subsidies and labor market behavior* undersøger jeg effekten af et øget befordringsfradrag på lønninger, arbejdsløshed, og afstand til arbejde for de berørte indbyggere. Selv om de fleste lande bruger store summer på subsidier til pendlere er virkningerne af tilskuddet ukendte. Dette kapitel bidrager til litteraturen ved ikke blot at undersøge selve befordringsfradraget, men også ved at udlede kausale effekter på nøgle faktorer på arbejdsmarkedet. Identifikationen stammer fra en dansk skattereform, hvor befordringsfradraget blev hævet i udvalgte udkantskommuner, men ikke i resten af landet. Ved brug af metoderne propensity score matching og difference-in-difference estimerer jeg den gennemsnitlige effekt på de berørte. Resultaterne viser, at beboerne i de områderne med mulighed for højere befordringsfradrag øger afstanden til arbejde, men arbejdsløsheden i de pågældende områder er upåvirket af det hævede fradrag. Flere reformer, der skulle hjælpe udkantsområder har vist sig mere eller mindre virkningsløse.

I det fjerde kapitel med titlen *The cost of bargaining: The case of fathers' parental leave* udforsker jeg teoretisk de kendte forhandlingsmodeller indenfor familien. Jeg kommer frem til at den bedste model til at beskrive fædres barsel er *The gendered spheres model.* I stedet for at parret går fra hinanden, når de ikke kan samarbejde, påtager de sig de opgaver, som er indenfor deres køns 'sphære'. Når vi taler om barsel, er denne sphære feminim, hvorfor barsel som udgangspunkt tilhører moren. Hun skal være villig til at afgive noget af barslen og faren skal være villig til at påtage sige noget af barslen for at fædre kan komme på barsel. Modellen kan være en forklaring på, hvorfor så få danske mænd er på barsel. Udover at de kønnede sphærer gør det svært at dele barslen, er høje forhandlingsomkostninger med til at gøre det endnu sværere. Hvis et par oplever store omkostninger forbundet med at diskutere arbejdsdeling, vil de kønnede sphærer i endnu højere grad betyde lavere barsel for mænd. Det finder jeg også empirisk. Undervejs sammenligner jeg det danske frie valg, hvor forældre selv afgør, hvordan de deler barsel, med kvote-systemet i resten af Skandinavien, som øremærker en del af barslen kun til faren. I den empiriske model, estimerer jeg ved hjælp af en lineær sandsynlighedsmodel sammenhængen mellem forhandlingsomkostninger og fædres barsel. I den ene af to delanalyser, finder jeg at forhandlingsomkostninger sænker fædres barsel mere i familier, hvor parret har lige meget magt eller hvor faren har mere magt, end i familier, hvor moren har mest magt. Her forstås magt som en stærk forhandlingsposition. I den anden delanalyse viser resultaterne at for forældre der allerede har børn, afskrækker høje forhandlingsomkostninger fars barsel mere end for førstegangs forældre. Jeg konkluderer at systemet svækker samarbejdet mellem forældre, hvis forældre alene beslutter deres deling af barselsorlov.

Education Economics, 2013
http://dx.doi.org/10.1080/09645292.2013.789825

Increasing the admission rate to upper secondary school: the case of lower secondary school student career guidance

Anders Hoest[a], Vibeke Myrup Jensen[a] and Lisbeth Palmhoej Nielsen[a,b]*

[a]*The Danish National Centre for Social Research, Copenhagen, Denmark;* [b]*Department of Economics and Business, Aarhus University, Aarhus, Denmark*

(*Received 14 August 2012; final version received 22 March 2013*)

Although several studies investigate the effects of school resources on student performance, these studies tend to focus more on intervention effect sizes than on their cost-effectiveness. Exploiting policy-induced variation in Denmark and using high-quality administrative data, we investigate the effects of a school intervention that introduces structured student career guidance in lower secondary school on upper secondary school admission. Disregarding the sunk-cost of implementation, the reform was cost-neutral. In a difference-in-difference framework, we find that the reform increases admission to upper secondary school between 4.0 and 6.3 percentage points for immigrants, but shows at best small improvements for the native students.

Keywords: policy evaluations; difference-in-differences; career guidance; upper secondary school

1. Introduction

Ageing populations and concerns about future tax bases are increasingly difficult to ignore in most industrialised countries. As raising the educational level is one means of increasing tax payments, several papers investigate the effects of school interventions, such as decreasing class size, increasing student−teacher ratio, and enhancing teacher quality (Angrist and Lavy 1999; Fredriksson, Ockert, and Oosterbeek 2013; Hanushek and Rivkin 2006; Krueger and Whitmore 2001). However, except for Fredriksson, Ockert, and Oosterbeek (2013), these studies tend to focus more on effect sizes than on the cost-effectiveness of the interventions. At a time where most countries are facing budget cuts, the question remains: how can we increase educational attainment without overspending? This paper evaluates a school reform in Denmark designed to increase the quality of student guidance. More specifically, we identify the effects of the intervention on admission to upper secondary school.

Introducing more structure, coherence, and quality into the guidance system, in 2004, the Danish Ministry of Education (DME) changed the guidance system in all public schools. The scope was twofold: to improve the admission rate to upper secondary school and reduce the upper secondary school dropout rate.[1] To accomplish this task, career counselling was centralised in regional centres, through which the municipalities expected to gain economies of scale from larger units and knowledge spill-overs between counsellors. Before the reform, school counselling practices were mainly carried out locally, with no national requirements.

*Corresponding author. Email: lpn@sfi.dk

Methodologically, this paper follows Machin and McNally's (2008) influential study 'The Literacy Hour', which investigates the effect of a small school intervention. Estimating a difference-in-differences (DiD) model, we use public schools as the treatment group, because the intervention was implemented only in public schools, and use private schools as the control group. The private schools follow the same curriculum standards as the public schools, and the parents pay only about 25% of the actual cost, with the rest funded by the municipalities. Therefore, we argue that only minor differences exist between Danish public and private schools.

Using high-quality administrative registers on school resources, with detailed information on students and their parents, we follow six cohorts of ninth graders (final year of compulsory schooling) into their first year of upper secondary school.

The literature on student guidance is sparse. Most earlier studies present descriptive measures instead of causal effects (see Hughes and Karp 2004 for an overview), investigate the effects of students' attitudes instead of their actual behaviour (McKay, Bright, and Pryor 2005), or focus on a subsample of the population (e.g. upper secondary school graduates) (Bettinger et al. 2012; Borghans, Golsteyn, and Stenberg 2011). This paper analyses the effect of the reform on complete cohorts.

The theory that students fail to enrol in the post-secondary education due to a lack of information about how to succeed has been proven wrong by analyses of information shocks and their effect on student take-up of loans (Booij, Leuven, and Oosterbeek 2012) and inclination to apply for post-secondary education funding (Bettinger and Baker 2011). Both studies find that the students acquired the information but failed to apply it. Combining information with individual guidance, however, is apparently proved highly effective (Bettinger et al. 2012). The Danish reform combines general information with in-class and individual student guidance.

The immigrants are of special interest because they have larger drop-out rates (Bratsberg, Raaum, and Røed 2011), they experience larger barriers to education through less information about the application process, and they have larger gains from schooling than the native-born students (Perna et al. 2008; Turney and Kao 2009), especially if they grow up amongst uneducated adults (Åslund et al. 2011). In Norway, Brinch, Bratsberg, and Raaum (2012) find large effects of a non targeted reform, such as the Danish Guidance Reform (hereafter DGR), especially on immigrants' school attainment. Thus, the immigrants are a vulnerable group that may have potentially high gains from improved career guidance and information.

We find that the new guidance structure improves admission to upper secondary school by 4.0–6.3 percentage points for the immigrants but no robust effect for the native-born, where we find sizes between no effect and 1.0 percentage points. We perform several sensitivity checks to validate our results, and in total we find that except for including school-specific time trends, our results for the immigrants are robust to these checks, whereas our results for the native-born are sensitive to small changes to the model specification.

2. Background

Primary or lower secondary school in Denmark consists of grade 0 (ages 5–6) to grade 9 (ages 15–16) and an optional grade 10. After grade 9 or 10, students either enter upper secondary school or leave the educational system. Students that do not directly enrol in upper secondary school may enter later, with no loss of rights or opportunities for enrolment. For the 2002 grade 9 cohort, 5.1% of the native-born and 10.1% of the immigrants had not enroled in upper secondary school with 5 years after leaving grade 9 (DME 2012). As the upper secondary school consists of several tracks (academic and vocational), career guidance aims at motivating students and helping them to choose between the tracks according to their abilities and wishes. Guidance is mainly carried out in the grades 8–10 in both public and private schools.

2.1 *The 2004 change in the student career guidance system*

In 2002, an OECD report pointed out three major weaknesses in the Danish guidance system: an inward-looking sectorial guidance system (lower secondary school, upper secondary school, etc.), with little continuity across sectors; a low educational level of guidance counsellors; and a lack of effective quality-assurance procedures (OECD 2002). On 1 August 2004, DME implemented the DGR addressing the weaknesses pointed out by the OECD (Jensen and Frederiksen 2004).

Before the DGR – at both public and private schools – the school principal was responsible for career guidance, and the classroom teachers, with a wide range of short-term training courses in counselling, executed the guidance activities. No national or regional requirements existed in terms of the content of the career guidance or the level of counselling qualifications. The municipalities are the local school authorities, and although before the reform the annual grants from the municipalities to the schools included funding for career guidance, the money was not directly earmarked for counselling.

With the DGR, the DME centralised the organisation of the career guidance, making new regional centres responsible for executing the career guidance at primary and upper secondary schools.[2] The guidance counsellors work full-time and are trained through a six-month full-time training programme (courses at the tertiary level). In cooperation with the municipalities, the centres determine the regional activities of the career guidance and, although no specific national requirements exist, the centres are required to document admission, drop-out rates, etc. After the reform, the municipalities remain responsible for administering the budget (DME 2004). Table 1 summarises the key features of the career guidance systems before and after the reform.

In total, the DGR constitutes a change from a *laissez-faire* system to a structured and highly professionalised setting focusing on more qualified guidance. As the centres also have the authority to reallocate their resources amongst the schools and the students, the new guidance targets mainly students at risk of leaving the education system after compulsory schooling or of dropping out of either lower or upper secondary school.

Table 1. A comparison of the student guidance system before and after the reform.

	Pre-reform period	Post-reform period
Management	School principal	Counsel centres
Student counsellors	Classroom teachers	Highly trained counsellors
Requirements	No requirements or follow-up	Centres are obligated to document admission and drop-out rates
Funding	School budgets	Centralised at the centres
Counsellors' training	20 different short-term training courses	One course, corresponding to six month full time university training

Source: Jensen and Frederiksen (2004).

Until 2007, the DGR did not encompass private schools. However, during 2005–2007, private schools could purchase career guidance at the regional guidance centres, or the private schools' counsellors could cooperate with the regional centres without additional costs (Association of Private Schools (APS) 2004). Nonetheless, neither the DME nor the APS has data on the number of private schools cooperating with their local student guidance centre before 2007.[3]

According to the DME Act No. 298 (2003), the DGR is cost-neutral, as the funds otherwise earmarked for career guidance at the lower and upper secondary schools finance the centres' operating cost. However, because municipalities chose different types of constellations of the new guidance system, the cost of that system also varied between municipalities. Thus, we use the guidance programme in Copenhagen, the capital of Denmark, to exemplify the start-up expenses of the DGR. In Copenhagen, the total budget for the first year was 232.48 euro per student (calculated on the basis of students in grade 8–10 in public schools and all students in upper secondary school). The sunk-costs on recruiting guidance counsellors, establishing offices and other facilities amounted to 4.76 euro per student (2% of total budget), and the government joined in with 19 euro per student (8% of the total budget) first year. All in all the start-up-expenses the first year amounted to 10% of the total spending per student (Copenhagen city council 2004). Generally, as the municipalities were responsible for reallocating resources to implementing the DGR, some municipalities may have financed sunk-costs by making the career guidance more efficient; others by cutting expenditures in other areas, such as public administration.

3. Empirical strategy

Given the natural experiment setting in which the DGR exposed some schools to changes in the career guidance system but not others, we use a DiD strategy in which public schools are the treatment group and private schools are the control group. Ninth-grade students in 2002–2004 and 2005–2007 define the pre- and post-reform periods, respectively. The standard DiD model for student *i* at school *s* in time *t* is the following:

$$\text{Admis}_{ist} = \beta_{0s} + \beta_1(\text{Public}_s{}^*\text{DGR}_t) + \beta_2\text{DGR}_t + \beta_3\text{Year}_t + X_{ist}\beta_4 + \varepsilon_{ist}, \quad (1)$$

where Admis_{ist} represents admission to upper secondary school. Public_s is a dummy variable equal to 1 for public schools and 0 for private schools. DGR_t is a dummy variable, where 1 defines the period with the new career guidance system and 0 the old system. The interaction between Public_s and DGR_t defines the variable of interest: the average treatment effect of the DGR on the treated.Year_t is a set of school year dummies capturing the time trend. X_{ist} represents a vector of individual, parental and school characteristics and ε_{ist} is the individual-specific error term. We estimate the model as a school-fixed effects model and thereby encompass any time-constant effects common to schools. Therefore, β_{0s} defines the school-specific constant and the empirical model the DGR_t dummy drops out. We use a linear probability model to estimate Equation (1), computing robust standard errors clustered by schools.

3.1 *Private schools as the control group*

The DGR constitutes a change from *laissez-faire* to structured guidance in public schools. If, and only if, private schools are a valid comparison group, we can identify any causal effects of structured student guidance on admission to upper secondary school. As only minor differences exist between public and private schools in Denmark, we argue that private schools are indeed a relevant control group for public schools.

Fifteen per cent of all students enrol in private schools, and DME defines the minimum educational standards for both public and private schools. Whilst public schools are fully funded through regional and state taxes, private schools receive funding equivalent to 75% of the cost of the average public school student. Furthermore, Rangvid (2008) finds that the difference between student achievement in private and public schools disappears when she controls for background characteristics. Consequently, public and private schools are alike in Denmark.

However, the *choice* of public or private school is endogenous, and changes in the composition of public and private school students across time could potentially contaminate the estimations (Blundell and Costa-Dias 2009). Some parents choose private schools because they wish their children to be more engaged in religion, music, sports, or languages; others, because they are unsatisfied with the local public school; and still others prefer private schools run entirely by a board of parents (Pedersen 2010). As the DGR received little attention in the daily press before it came into force, we find it highly unlikely that private school students would change to public schools due to the reform. Figure 1 shows the number of articles every month

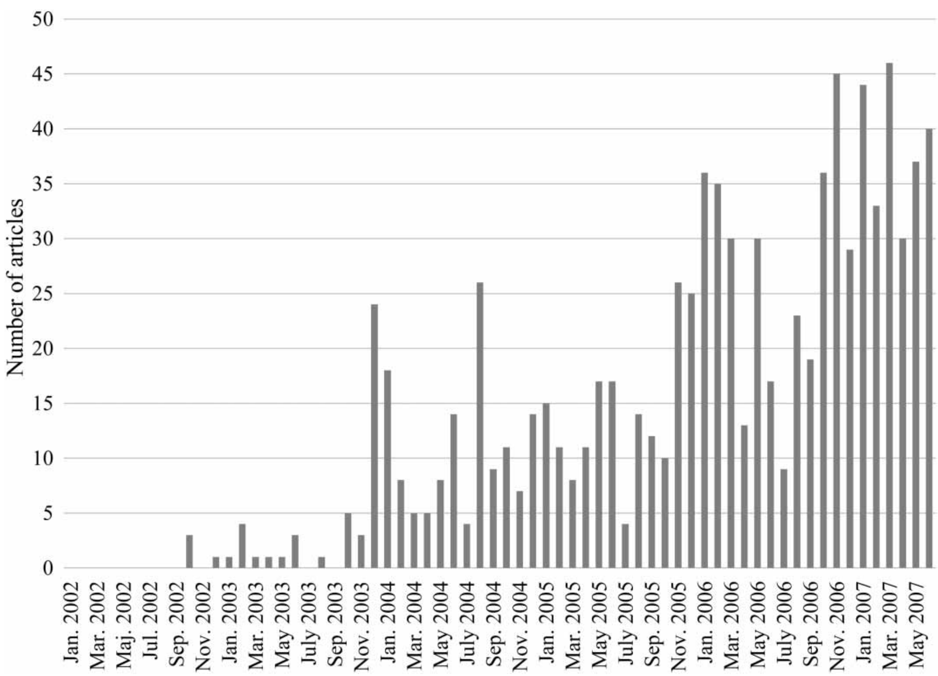

Figure 1. Number of articles mentioning the DGR in the pre- and post-reform period (January 2002–June 2007).
Notes: The figure counts the number of articles mentioning the DGR in national or local newspapers per month from January 2002 to June 2007. We use the following search criteria [in Danish] 'UU-vejledning', 'Ungdommens Uddannelses Vejledning', or 'Ungdommens Uddannelsesvejledning'. Source: www.Infomedia.dk

mentioning the DGR from January 2002 to June 2007. From January 2002 until July 2004 – before the reform – the DGR is mentioned about 100 times in the local and national newspapers. From August 2004 and onwards, the number of articles mentioning the new guidance system rises steadily predominantly in the local newspapers yet is still mentioned only about 80 times in the national newspapers.

Furthermore, we also find that for all four groups (native-born and immigrants in public and private schools, respectively), the percentage of students enroled at the same school in grades 8 and 9 is constant around the implementation of the reform. Therefore, we conclude that the DGR did not affect enrolment patterns one year after the reform (see appendix Figure A1 in the working paper version of this paper: Hoest, Jensen, and Nielsen 2012).[4]

Although we do not believe that parents changed schools because of the DGR (as the reform was a minor change), we cannot rule out that some parents did. In Section 4.2, we, therefore, investigate the student composition and enrolment changes concurrent with the reform.

4. Data

4.1 *The samples*

The data contain most ninth-grade students enroled at a public or a private lower secondary school during the school years 2002–2007 and combines various high-quality administrative registers from Statistics Denmark and DME. Information on school type and school resources stems from DME and through unique institution and individual identifiers, we merge these school-level records to individual-level information on the students and their parents. We observe all students in the ninth grade and the following two years. Thus, the data have a panel structure at the school level but takes the form of repeated cross sections at the individual level.

The outcome variable is a dummy identifying admission to the academic or the vocational track of upper secondary school. As about half of the 9th grade students attend 10th grade, we define the outcome as equal to 1 if the students enrol in upper secondary school in the first or second year after completing 9th grade, and otherwise 0. As some secondary educations take up to four years to complete, we are limited to examining admission rates rather than completion. However, the literature also finds that beginning sooner rather than later also increases completion. Dobkin and Ferreira (2010) find that younger students, despite lower achievement than their classmates, complete high school at higher rates than their older peers. Furthermore, Bailey, Jeong, and Cho (2010) report that older students are more likely to drop out of post-secondary education. If the DGR increased the admission rate, these findings indicate that the students analysed in this study may also have a higher likelihood of completion.

DME (2012) calculates the expected completion rates of upper secondary school five years after ninth grade. As one purpose of the DGR is to allocate extra guidance resources to low-performing students, we analyse not only native-born students (69% completion rate) but also immigrants (55% completion rate). We follow Statistics Denmark in defining 'immigrants' as both first- and second-generation immigrants, i.e. students who were born abroad and whose parents were born abroad (first generation) and parents are born abroad but students born in Denmark (second generation). Students with at least one parent born in Denmark, irrespective of origin at birth, are considered native-born.

Our final data set consists of 210,546 grade 9 native students enroled in 998 public (168 private) schools, and 15,013 grade 9 immigrant students enroled in 262 public (23 private) schools. These samples are equivalent to 77% of all public schools; however, as many private schools offer schooling only to grade 6 or do not offer 9th-grade achievement testing, the samples contain 40% of all private schools. For both the native-born and the immigrants, we restrict the data to schools that for all 6 years have at least 3 grade 9 students, aged 14–17, and a minimum of 1 grade 9 student achievement test in the 2 main subjects: maths and Danish.[5] In addition, as some schools have no immigrants, we restrict schools in the immigrant sample to having a minimum of three immigrants annually. This restriction ensures that we calculate the effects of DGR on a minimum of 18 immigrant students per school.

For the native-born, Figure 2 shows the distribution of 9th-grade enrolment in private and public schools, and illustrates that most (>50%) public schools have 15–30 grade 9 students per year, whereas most private schools have 5–18 students in grade 9.

Similarly, for immigrants most public and private schools in our sample have only 3–5 immigrants in grade 9 (Figure 3). In addition, demanding at least three immigrant students per school, per year, we exclude about 50% of all the schools that have immigrants.[6]

Figure 4 illustrates the annual admission rate to upper secondary school from 1982 to 2007 for the treatment and control groups, where the vertical line between 2004 and 2005 defines the implementation of the DGR. The figure shows that students from public schools (the solid lines) have a lower admission rate than students from private schools (the dotted lines). The dotted vertical line in 2002 shows the beginning of our period.

The native-born students in public schools have an almost constant admission rate around 72% throughout the period. For the native students in private schools, the

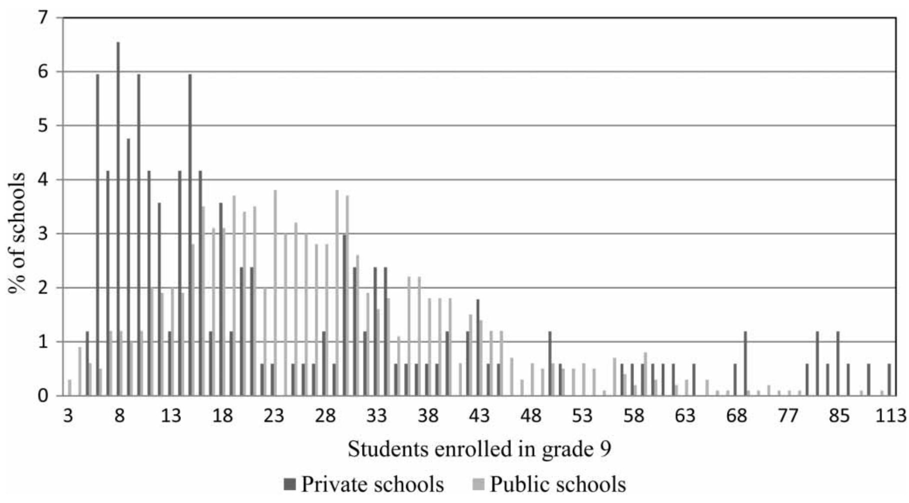

Figure 2. The distribution of grade 9 enrolment for the native-born students, by public and private schools.
Notes: The figure illustrates the distribution of grade 9 enrolments for the native-born students by private and public schools. The vertical axis defines the percentage of schools and the horizontal axis defines the number of students enroled.

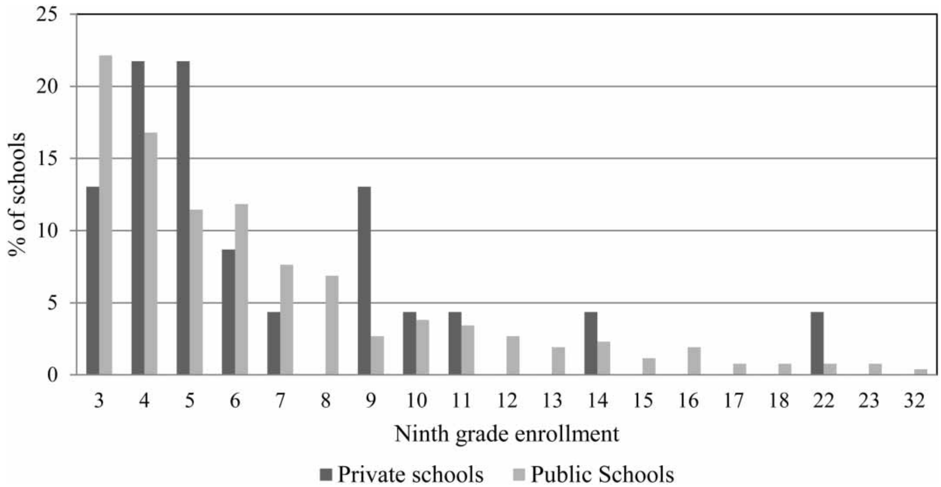

Figure 3. The distribution of grade 9 enrolment for immigrants, by public and private schools.
Notes: The figure illustrates the distribution of grade 9 enrolments for immigrant students by private and public schools. The vertical axis defines the percentage of schools and the horizontal axis defines the number of students enroled.

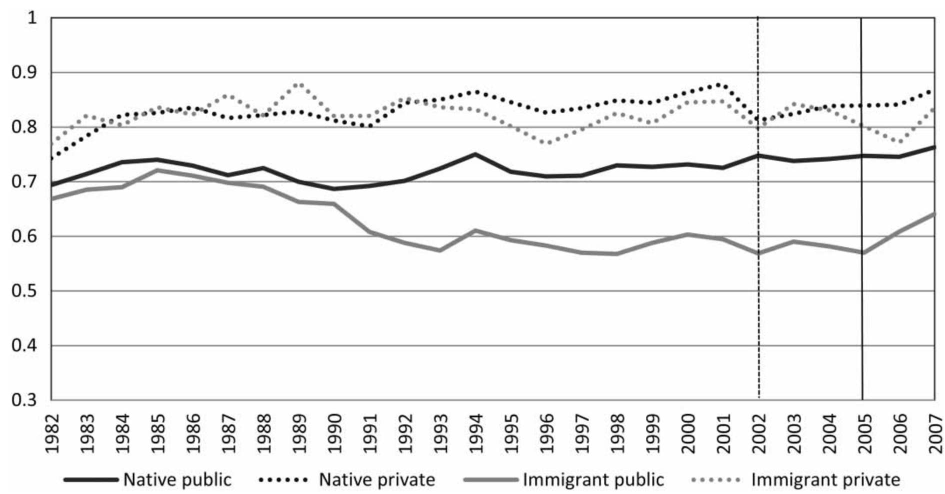

Figure 4. Admission to upper secondary school by native-born and immigrant students in public and private schools, respectively.
Notes: For each grade 9 cohort, the figure shows the proportion of students entering upper secondary school. The vertical dotted line in 2002 defines the start of our data window, whereas the vertical solid line at 2005 symbolises the first cohort after the reform. Over the entire period, we measure the same schools, however, the data before 2002 are slightly less restrictive than after 2002. The sample from 2002 to 2007 equals our main sample in the data where we make the following restrictions: Schools need to have at least three students per year that have a minimum of one exam in maths or Danish, and all schools need to be in the sample in all six years. While we use the same sample of schools in the 1982–2001 periods as the 2002–2007 periods, we cannot be sure that these schools have at least three students taking an exam in maths or Danish, because no data on exams exist before 2002. In addition, all schools are not included in all years from 1982 to 2007, because some schools started later.

admission rate fluctuates around 85% until the reform and afterwards continues at a high level.

The immigrant students in public schools have a stable admission rate around 60% from 1994 and until the reform. After the reform, the admission starts to increase. For the immigrants in private schools, the admission rate fluctuates around 82%, decreases from 82% in 2003 to 79% in 2006, and then increases again in 2007. In the main analysis, we average over three pre-reform and post-reform years to minimise the effects of year-by-year fluctuations. Nonetheless, a dip in the control group is critical for our analysis: in Section 5.1, we test whether our results are sensitive to excluding 2006 and in Section 5.3 test for differential trends.

4.2 *The covariates composition*

As discussed in Section 3.1, endogenous changes in student composition around the time of the reform may potentially contaminate our results. Therefore, to substantiate a causal interpretation of the results, we investigate whether student characteristics change concurrent with the reform. Table 2 shows means and standard deviations for each covariate for the treatment group (columns 1 and 2) and the control group (columns 3 and 4) before and after the reform, respectively. The means show that the native-born select positively into private schools (e.g. better educated, higher-income parents).

Column 5 in Table 2 presents the unconditional DiD estimator – the trend in public school covariates minus the trend in private school covariates.[7] This estimator shows that despite the differences in levels, most covariates change only marginally (1–3 percentage points) with the reform and only some of these changes are significant. The differences are mainly due to the large sample size in which even little differences are likely to be significant. The most prevalent difference is in grade point average (GPA), where students in private schools have seen an increase in grades from before to after the reform. As GPA is an important indicator for student preparedness for upper secondary school, these strong effects of GPA undermine our parallel trend assumption and are likely to affect our results. However, the effect of GPA is negative, indicating that private schools are better at obtaining high achievements. Thus, we expect that excluding GPA will downward bias our estimates. In Table 3, we provide further evidence of how GPA affects our results.

Similar to Table 2, we estimate summary statistics and the unconditional DiD estimator for the immigrants. For this group, we only find that the group of children, where one of the parents is living with a new partner decreases significantly more for the children in private schools than children in public schools around the time of the reform (significant at the 5% level). For all other covariates, we find no significant differences at the 5% level (see appendix Table A1 in the working paper version of this paper: Hoest, Jensen, and Nielsen 2012). We thus conclude that public and private schools are comparable for the immigrant sample.

5. Results and robustness

5.1 *Main results*

This section presents the effects of the DGR on admission to upper secondary school for native-born and immigrants. In Table 3, columns 1 and 2 display the effects for

Table 2. Summary statistics and the unconditional DiD estimator for the native-born students.

	Public schools		Private schools		
	Pre-reform	Post-reform	Pre-reform	Post-reform	Unconditional DiD estimator
GPA					
A or B	0.02	0.03	0.04	0.05	−0.01**
C	0.51	0.49	0.59	0.60	−0.03***
D	0.28	0.27	0.25	0.26	−0.01*
E or Failed	0.15	0.18	0.09	0.09	0.03***
GPA missing	0.04	0.03	0.03	0.01	0.01**
Student, gender	0.51	0.51	0.47	0.47	0.01*
Student, age	15.49	15.58	15.46	15.54	0.00
Mother's education					
Primary or lower secondary schooling	0.26	0.22	0.19	0.14	0.00
Upper secondary school, academic track	0.04	0.05	0.05	0.06	0.00
Upper secondary school, vocational track	0.37	0.39	0.32	0.34	0.01
Tertiary education	0.31	0.32	0.42	0.45	−0.01**
Missing education	0.01	0.01	0.01	0.01	0.00*
Father's education					
Primary or lower secondary schooling	0.23	0.22	0.16	0.14	0.01***
Upper secondary school academic track	0.04	0.04	0.06	0.06	0.00
Upper secondary school, vocational track	0.42	0.42	0.35	0.36	−0.01
Tertiary education	0.26	0.27	0.37	0.38	0.00
Missing education	0.03	0.03	0.03	0.03	0.00
Mother's employment					
Permanent social benefits	0.04	0.04	0.04	0.03	0.00**
Temporary social benefits	0.08	0.06	0.06	0.05	0.00*
Missing employment	0.02	0.02	0.02	0.02	0.00
Father's employment					
Permanent social benefits	0.04	0.04	0.03	0.03	0.00
Temporary social benefits	0.05	0.04	0.04	0.03	0.00
Missing employment	0.02	0.02	0.03	0.03	0.00
Mother's income quartile					
Income Q1 – lowest	0.22	0.22	0.19	0.17	0.02***
Income Q2	0.25	0.25	0.20	0.20	0.00
Income Q3	0.25	0.25	0.23	0.25	−0.02***
Income Q4 – highest	0.24	0.24	0.33	0.34	−0.01
Missing income	0.04	0.04	0.05	0.04	0.00
Father's income quartile					

(*Continued.*)

Table 2. (*Continued.*)

	Public schools		Private schools		
	Pre-reform	Post-reform	Pre-reform	Post-reform	Unconditional DiD estimator
Income Q1 – lowest	0.21	0.20	0.18	0.16	0.02***
Income Q2	0.25	0.25	0.18	0.19	0.00
Income Q3	0.25	0.25	0.23	0.24	−0.01*
Income Q4 – highest	0.24	0.24	0.35	0.35	0.00
Missing income	0.05	0.05	0.06	0.06	0.00
Mother is unknown	0.00	0.00	0.00	0.00	0.00
Father is unknown	0.01	0.01	0.01	0.01	0.00
Mother is deceased	0.01	0.01	0.01	0.01	0.00
Father is deceased	0.02	0.02	0.02	0.02	0.00
Family Structure					
Nuclear	0.67	0.65	0.68	0.68	−0.02***
Divorced with new partner	0.12	0.13	0.10	0.11	0.00
Single parent	0.16	0.18	0.16	0.17	0.01**
Student living alone	0.02	0.02	0.03	0.01	0.01*
Missing family structure	0.02	0.02	0.03	0.03	0.00
Number of siblings					
One sibling	0.50	0.50	0.50	0.51	−0.01
Two or more siblings	0.45	0.45	0.43	0.42	0.01
Age of mother at time of birth					
Mother's age missing	0.01	0.01	0.01	0.01	0.00
Teenage mother	0.02	0.02	0.01	0.01	0.00
Student, health at birth					
Admitted to hospital up to three years after birth	0.94	0.93	0.92	0.92	0.00
Low birth weight (<2500 grams)	0.05	0.05	0.04	0.04	0.00
Missing birth weight	0.02	0.02	0.03	0.03	0.00
Premature birth (<37th week)	0.08	0.09	0.08	0.08	0.00
Missing, gestation	0.02	0.02	0.03	0.03	0.00*
Number of observations	106,455	116,394	15,723	16,853	255,425

Notes: The unconditional DiD estimators are calculated according to model (2), i.e. the differences in the means between the treatment and the control group, before and after the intervention. We apply school-level fixed effects when identifying the level of significance.

*Significance at the 10% level.

**Significance at the 5% level.

***Significance at the 1% level.

Table 3. The effects of the DGR on admission to upper secondary school.

	Native-born sample		Immigrant sample	
	(1) No covariates	(2) All covariates	(3) No covariates	(4) All covariates
Effect of guidance	−0.0168***	0.0098**	0.0478**	0.0626***
Standard errors	0.0053	0.0045	0.0198	0.0237
Number of observations	255,425	255,425	18,235	18,235
2006 excluded	−0.0173***	0.0064	0.0277	0.0528**
Standard errors	0.0056	0.0050	0.0205	0.0249
Number of observations	210,546	210,546	15,013	15,013
GPA excluded	–	−0.0066	–	0.0395**
Standard errors	–	0.0047	–	0.0189
Number of observations	255,425	255,425	18,235	18,235

Notes: For both the native and the immigrant sample, we use model (1) to calculate the estimates. Columns 1 and 3 represent estimations including year dummies as covariates. Columns 2 and 4 represent estimations including all covariates in Table 2.

*Significance at the 10% level.

**Significance at the 5% level.

***Significance at the 1% level.

the native-born, and columns 3 and 4 show the effects for the immigrants. Whilst columns 1 and 3 present results including only year dummies (hereafter, 'results with no controls'), columns 2 and 4 present results with the full set of controls (see Table 2 for the list of controls).

The first row of estimates in Table 3 shows the main effects. For the native-born, in the model with no controls, we find that the DGR decreases the probability of admission to upper secondary school by 1.7 percentage points. However, including all controls, the probability of admission to upper secondary school increases by 1 percentage point. For the immigrants, in the model with no controls, the DGR increases the probability of admission to upper secondary school by 4.8 percentage points. Including all controls increases the effect to 6.3 percentage points.[8] This result corresponds with a recent study by Borghans, Golsteyn, and Stenberg (2011), who find that career guidance in the Netherlands has the largest effect for immigrants.[9]

The second row of estimates in Table 3 presents results excluding data from 2006, due to the dip in the admission rate for private school students (Figure 4). Although this dip appeared mainly for immigrant admission, when we exclude 2006, the significant positive effect of the DGR disappears for the native-born. The effect for the native-born was small to begin with, and we believe that the effect disappears due to the reduction in sample size. Excluding 2006 from the immigrant sample, the effect remains significant but marginally smaller (5.3 percentage points). This lower effect suggests that our main effect is partly upward biased by the reduction in the admission rate for the untreated.

The third row of estimates in Table 3 presents results, where we exclude GPA. Table 2 showed a significant negative trend in GPA from before to after the reform.

The negative trend means that public schools have increased their grades less over the period than the private schools. To test the hypothesis that excluding GPA is likely to downward bias our results, we exclude GPA from our main specification and, as predicted, the estimates become smaller; for the native-born, even insignificant and negative. For the immigrants, however, the DGR still increases the admission rate to upper secondary schooling significantly by 4.0 percentage points.

The effects for the immigrants could be generated by students with the abilities to attend upper secondary school but who before the reform would meet other barriers to entrance. The immigrant students may have lacked not only information about upper secondary schooling, but also guidance in how to apply that knowledge. According to Bound, Lowenheim, and Turner (2010), admission to higher education in the USA has increased rapidly but completion less so. The authors argue that the lack of attainment is due to a lower quality of students entering higher education, a possible explanation for why we find no effects for the native students.

In the next three sections, we continue to investigate how sensitive our results are to our model specification.

5.2 *Treatment of the untreated*

Section 2.1 explained that private schools can enrol in the new student career guidance system if they pay an additional fee. Despite this opportunity, we found only few examples of private schools collaborating with the local counselling centre. Assuming that our results are only weakly biased by this opportunity, we proceed by investigating how sensitive our results are to this type of potential treatment of the untreated. We thus reallocate increasingly more of the private school students to treatment and then plot our new coefficients of interest and their corresponding 99% confidence intervals (CIs). For example, we first assume that 1% of private school students were treated. We then randomly sample 1% of private school students, recode them to treatment, and, using model 1, estimate the effect of the DGR on admission to upper secondary school. Then, we continue to a 2% random allocation of treatment, etc., until 50% of the private schools students are assumed treated. For each level of reallocation, we make 100 replications of each estimate and then compute the means and the CIs of these estimates.

On the basis of these estimations, we find that 30% of the native-born in private school can be treated without the effect of the DGR disappearing, whereas 35% of the immigrants can be treated. Second, for these threshold values, the effect of the DGR is 0.7 for the native-born and 4.3 percentage points for the immigrants (see appendix Figures A2 and A3 in the working paper version of this paper: Hoest, Jensen, and Nielsen 2012).

5.3 *Test for parallel time trends*

Parallel time trends in the treatment and control groups are an important assumption for the validity of the DiD estimator. We test this assumption in two ways: first by estimating a traditional falsification test and, second, by including school-specific trends into model (1).

First, we estimate the traditional pseudo-reform falsification test. If we assume that potential differences in the time trends between the treatment and control groups are constant, we can use a DiD model similar to model (1) using only data before 2005 to test for the effects of differential trends between the two groups. We apply this

Table 4. Pseudo-reform falsification test: the effect of the DGR on admission to upper secondary school using either 2003 or 2004 as the false post-reform year.

	Native-born sample		Immigrant sample	
	(1) Pre-reform: 2002, Post-reform: 2003	(2) Pre-reform: 2002, Post-reform: 2004	(3) Pre-reform: 2002, Post-reform: 2003	(4) Pre-reform: 2002, Post-reform: 2004
Pseudo-reform effect	0.0042	0.0006	−0.0265	−0.0201
Standard errors	0.0079	0.0088	0.0455	0.0378
Number of students	80,290	82,167	5427	5577
Number of schools	1166	1166	285	285
R^2	0.335	0.338	0.406	0.418

Notes: Columns 1 and 3 use 2002 as the pre-reform year and 2003 as the false post-reform year. Columns 2 and 4 use 2002 as the pre-reform year and 2004 as the false post-reform year. For both the native-born and the immigrants, we use model (1) and include all covariates presented in Table 2.

*Significance at the 10% level.

**Significance at the 5% level.

***Significance at the 1% level.

robustness check using students from 2002 as the pre-reform period and students from 2003 or 2004 as the false post-reform periods.

Table 4 presents the results of the falsification test. For native-born, column 1 shows the effects using 2003 as the false post-reform period and column 2, the effects of using 2004. For the immigrants, columns 3 and 4 show the equivalent estimated effects. For the identification strategy to be valid, this falsification test should provide small effects, as changes in admission outside the reform are endogenous. For the native-born, we find close to zero and insignificant effects. For the immigrants, we also find insignificant effects, but the point estimates suggest that the false reform decreases admission to upper secondary school by 2.7 percentage points and by 2 percentage points when we use 2003 and 2004 as the false post-reform period, respectively. Whilst we do not find zero effects of this falsification test for immigrants, the negative results suggest that this trend is not upward-biasing our results.

Second, we test the model for its sensitivity to school-specific time trends. We find no significant effects of the DGR for either the native-born or the immigrants. Thus, the results suggest that school-specific time trends are potentially generating our main results. However, including school-specific time trends is a very strict test, as in principle it means that a specific trend on the basis of three students per year is calculated for some schools (see appendix Table A4 in the working paper version of this paper: Hoest, Jensen, and Nielsen 2012).

6. Conclusion

This paper estimates the effects of student career guidance in lower secondary school on admission to upper secondary school. We exploit a national change, in Denmark, in structure and quality in lower secondary school student career guidance, a change that affected students in public but not in private schools.

For both the native-born and the immigrants, we find that the DGR has a positive significant effect on admission to upper secondary school. However, for the native-born, the result is small (no effect to 1 percentage point), whereas for the immigrants the effect is larger (4.0–6.3 percentage points).

We perform several sensitivity checks. First, we find only minor changes in the covariates concurrent with the reform, except for GPA. Because GPA changes negatively with the reform, we estimate the model both with and without GPA. Second, we find no changes in the enrolment pattern into private and public schools around the reform. Third, a circular dip in the admission rate for immigrants in private schools exists around 2006, so we expand the analysis to include a model without 2006. Fourth, we test how sensitive our results are to potential treatment of the untreated. Fifth, we estimate how sensitive our results are to changes in time trend through a classic pre-reform falsification test and by including school-specific time trends. From these falsification tests, we conclude that our result for the immigrants is fairly robust, except for the inclusion of school-specific time trends. For the native-born, the result is sensitive to even small changes in the model specification.

The reform was designed to reallocate resources amongst the students so that mainly those with few resources received face-to-face guidance. As we find that mainly the immigrants benefit from the reform and that the new system, for the average native student, was as good as the old one, our results suggest that the change in the student career guidance system has had the expected impact.

As we cannot investigate the second main purpose of the new student career guidance system – the effect on upper secondary school dropouts – our study examines only half of the intended effects of the structured career guidance system. That effects on drop-outs can be found for the average native student is likely.

Acknowledgements

This paper arises from a report undertaken for the Danish Ministry of Education, by Jensen and Nielsen (2010). The article comes with a web appendix that will be published at the authors' web page after the article is accepted for publication. The authors acknowledge financial support by the Danish Agency for Science, Technology, and Innovation through grant number 09-065167 and by Aarhus University grant number 11-743-065. We thank Paul Bingley, Anders Holm, Dean Lillard, Giuseppe Migali, Marianne Simonsen, participants of Strategic Research Council Workshop at SFI 2012, participants of the ESPE 2012 conference, participants of the IWAEE 2012 and two anonymous referees for helpful comments. We also thank Natalie Reid for excellent writing assistance.

Notes

1. We cannot investigate the effect on the dropout rate, because this analysis requires data from a minimum of four years after the reform – and three years after the reform the student career guidance system changed again to allow all types of students to benefit.
2. Some centres only administer schools within one municipality, whereas others include schools from several municipalities. Thus, the municipalities could decide on different solution frameworks: they could establish new institutions alone or in cooperation with other municipalities, delegate responsibility for the career guidance to established public institutions or to independent institutions or private companies through invitations to tender. As part of the DGR, the DME established a national virtual career guidance portal containing information about upper secondary schools. This portal is publicly available for students

in both public and private schools and aims at helping students and their parents find relevant information about upper secondary schools. However, several websites provided a similar service throughout the period, and the portal is of minor interest relative to the effect of the DGR.
3. Legally, the municipalities are allowed to change the grants to private schools only three years after the change has been implemented in public schools.
4. Estimating the potential effect of the DGR on enrolment two years after the reform requires information on changes in enrolment in grade 7. Unfortunately, in our period students are not observed until grade 8.
5. Schools with fewer than 5 grade 9 students live predominantly in rural areas. The age limit excludes adult learners.
6. However, we find that relaxing this restriction and including schools with only one or two immigrants per year, the effects of the DGR are similar in magnitude. Excluding schools with a minimum of four through seven immigrants, the point estimate remains but loses precision presumably from the reduction in sample size (see appendix Table A3 in the working paper version of this paper: Hoest, Jensen, and Nielsen 2012).
7. For each covariate, the following model calculates the unconditional DiD estimator:

$$\text{Unc\_DiD}_s = (\text{public}_{s,t} - \text{public}_{s,t-1}) - (\text{private}_{s,t} - \text{private}_{s,t-1}), \quad (2)$$

where public_s defines the public schools and private_s defines the private schools before $(t-1)$ and after (t) the reform.
8. Whilst our preferred model is a linear probability model including school fixed-effects, we also try to estimate the effect using a probit model. Comparing the estimates from the linear probability model (without fixed effects) with the marginal effects from the probit model, we find similar effects of the DGR reform for the immigrants but somewhat larger effects in the probit model for the native-born (see appendix Table A2 in the working paper version of this paper: Hoest, Jensen, and Nielsen 2012).
9. We also investigate potential heterogeneous effects of the reform amongst students with different parental backgrounds by estimating triple-difference models. We find no significant effects and thus no not report these results.

References

Angrist, J., and V. Lavy. 1999. "Using Maimonides' Rule to Estimate the Effect of Class Size on Scholastic Achievement." *Quarterly Journal of Economics* 114 (2): 533–575.

Åslund, O., P.-A. Edin, P. Fredriksson, and H. Gronqvist. 2011. "Peers, Neighborhoods, and Immigrant Student Achievement: Evidence from a Placement Policy." *American Economic Journal: Applied Economics* 3 (2): 67–95.

Association of Private Schools (APS). 2004. *Nye regler for vejledning* [New Rules for the Career Guidance]. New about Private Schools. No. 5, October 2004. Copenhagen: Association of Private Schools.

Bailey, T., D. W. Jeong, and S. W. Cho. 2010. "Referral, Enrollment, and Completion in Developmental Education Sequences in Community Colleges." *Economics of Education Review* 29 (2): 255–270.

Bettinger, E., and R. Baker. 2011. "The Effects of Student Coaching in College: An Evaluation of a Randomized Experiment in Student Mentoring." NBER Working paper no. 16881. National Bureau of Economic Research at Cambridge, MA.

Bettinger, E., B. T. Long, P. Oreopaulos, and L. Sanbonmatsu. 2012. "The Role of Simplification and Information in College Decisions: Results from the H&R Block FAFSA Experiment." *The Quarterly Journal of Economics* 127 (3): 1205–1242.

Blundell, R., and M. Costa-Dias. 2009. "Alternative Approaches to Evaluation in Empirical Microeconomics." *Journal of Human Resources* 44 (3): 565–640.

Booij, A. S., E. Leuven, and H. Oosterbeek. 2012. "The Role of Information in the Take-Up of Student Loans." *Economics of Education Review* 31 (1): 33–44.

Borghans, L., B. H. H. Golsteyn, and A. Stenberg. 2011. "Does Expert Advice Improve Educational Choice." Paper presented at the EALE meeting, Cyprus, September 22–24.

Bound, J., M. F. Lowenheim, and S. Turner. 2010. "Why Have College Completion Rates Declined? An Analysis of Changing Student Preparation and Collegiate Resources." *American Economic Journal, Applied Economics* 2 (3): 129–157.

Bratsberg, B., O. Raaum, and K. Røed. 2011. "Educating Children of Immigrants: Closing the Gap in Norwegian schools." Discussion Paper series, Forschungsinstitut zur Zukunft der Arbeit, No. 6138, IZA, Bonn.

Brinch, C. N., B. Bratsberg, and O. Raaum. 2012. "The Effects of an Upper Secondary Education Reform on the Attainment of Immigrant Youth." *Education Economics* 20 (5): 447–473.

Copenhagen city council. 2004. "Dagsorden for Ordinært møde torsdag den 25." Marts 2004 [Meeting agenda on March 25, 2004]. Municipality of Copenhagen.

Danish Ministry of Education. 2004. *Guidance Education – A New Guidance System in Denmark*. Copenhagen: Reproff Print.

Danish Ministry of Education. 2012. *Profilfigurer* [The Profile Model]. The Danish Ministry of Education. Accessed February 15. http://www.uvm.dk/Service/Statistik/Tvaergaaende-statistik/ Andel-af-en-aargang-der-forventes-at-faa-en-uddannelse/Profilfigurer.

Danish Ministry of Education Act No. 298. 2003. *Lov om vejledning om valg af uddannelse og erhverv* [Act on Career Guidance – Education and Occupation]. The Danish Ministry of Education Act No. 289, April 30, 2003.

Dobkin, C., and F. Ferreira. 2010. "Do School Entry Laws Affect Educational Attainment and Labour Market Outcomes?" *Economics of Education Review* 29 (1): 40–54.

Fredriksson, P., B. Ockert, and H. Oosterbeek. 2013. "Long-Term Effects of Class Size." *The Quarterly Journal of Economics* 128 (1): 249–285.

Hanushek, E. A., and S. G. Rivkin. 2006. "Teacher Quality." In *Handbook of the Economics of Education*, Vol. 2, 1051–1078. Amsterdam: North Holland.

Hoest, A., V. M. Jensen, and L. P. Nielsen. 2012. "Increasing the Admission Rate to Upper Secondary School: The Case of Lower Secondary School Student Career Guidance." SFI Working Paper series no. 03:2012. Copenhagen: The Danish National Centre for Social Research.

Hughes, K. L., and M. M. Karp. 2004. *School-Based Career Development: A Synthesis of the Literature*. New York: Columbia University, Institute on Education and the Economy, National Training support Centre.

Jensen, S., and P. B. Frederiksen. 2004. *Reform af uddannelses- og Erhvervsvejledningen* [Reform of Educational and Occupational Guidance]. Uddannelse, February 2, 2004.

Jensen, V. M., and L. P. Nielsen. 2010. *Veje gennem ungdomsuddannelse 1, Statistiske analyser af folkeskolens betydning for unges påbegyndelse og gennemførelse af en ungdomsuddannelse* [Paths Through Upper Secondary Education. A Statistical Analysis of the Importance of Primary and Lower Secondary School for Entering and Completing Upper Secondary Education]. Copenhagen: The Danish National Centre for Social Research, Report 2010 10:24.

Krueger, A., and D. Whitmore. 2001. "The Effect of Attending a Small Class in the Early Grades on College-test Taking and Middle School Test Results: Evidence from Project Star." *Economic Journal* 111 (468): 34–63.

Machin, S., and S. McNally. 2008. "The Literacy Hour." *Journal of Public Economics* 92 (56): 1441–1462.

McKay, H., J. E. H. Bright, and R. G. L. Pryor. 2005. "Finding Order and Direction from Chaos: A Comparison of Career Counseling and Trait Matching Counseling." *Journal of Employment Counseling* 42 (3): 98–112.

OECD. 2002. *OECD Review of Career Guidance Policies.* Denmark – Country Note. Accessed April 2002. http://www.oecd.org/denmark/2088292.pdf.

Pedersen, L. K. 2010. *Flere vælger privatskole frem for folkeskole* [More Parents Choose Private Schools over Public Schools]. Folkeskolen. Accessed August 8. http://www.folkeskolen.dk/63340/flere-vaelger-privatskole-frem-for-folkeskole.

Perna, L. W., H. T. Rowan-Kenyon, S. L. Thomas, A. Bell, R. Anderson, and C. Li. 2008. "The Role of College Counselling in Shaping College Opportunity: Variations Across High Schools." *The Review of Higher Education* 31 (2): 131–159.

Rangvid, B. S. 2008. "Private School Diversity in Denmark's National Voucher System." *Scandinavian Journal of Educational Research* 52 (4): 331–54.

Turney, K., and G. Kao. 2009. "Barriers to School Involvement: Are Immigrant Parents Disadvantaged? *The Journal of Educational Research* 102 (4): 257–271.

The Effects of Maternal Employment on Children's Academic Performance

Rachel Dunifon, *Cornell University*
red26@cornell.edu

Anne Toft Hansen, *Danish National Center for Social Research*

Sean Nicholson, *Cornell University and NBER*

Lisbeth Palmhøj Nielsen, *Danish National Center for Social Research and Aarhus University*

Marts 2016

Abstract: Using data following over 335.000 Danish children from birth through 9th grade, we examine the effect of maternal employment during a child's first three and first 15 years on that child's grade point average in 9th grade. We address the endogeneity of employment in several ways: by including a rich set of household control variables, by instrumenting for employment with the gender- and education-specific local unemployment rate, and by including maternal fixed effects. Across most models, we find positive associations between maternal employment and children's academic performance. This is in contrast with the larger literature on maternal employment, much of which takes place in other contexts, and which finds no or a small negative effect of maternal employment on children's cognitive development.

Key words: maternal employment; educational outcomes; Denmark

Acknowledgements: Helpful comments have been provided by Kevin Milligan, Joseph Doyle, Marianne Simonsen, Nabanita Datta Gupta, participants at the ESPE conference, and seminar participants at Cornell University, Aarhus University, AU RECEIV center (project no. 908792), and the Danish National Center for Social Research (SFI).

I. Introduction

The labor force participation rate of mothers with young children has increased substantially in developed countries over the past 50 years. In the United States, for example, 71% of mothers were working in 2012 (Bureau of Labor Statistics, 2014). Given the continued primacy of mothers as the main providers of child care (Bianchi 2000; Bianchi, Milkie, Sayer, and Robinson, 2000), increases in maternal employment lead to concerns about the tradeoffs working mothers must face in terms of time investments in work vs. time spent at home. In the United States, research suggests that early and intense maternal employment may be linked to small declines in children's cognitive development (Hill, Waldfogel, Brooks-Gunn, and Han, 2005; Brooks-Gunn, Han, and Waldfogel, 2002; Ruhm, 2008). In Scandinavia, however, the associations between early maternal employment, and the child's cognitive development seem to be missing (Rasmussen, 2010; Liu and Skans, 2010; Dahl, Løken, Mogstad and Salvanes, 2013).

Heckman and Cunha (2007) model the skill-formation of children by using early child-investments as complementary building blocks for later investments. They argue that if early investments are not followed up by later investments, the gains from the early investment will be lessened. The empirical literature is mixed as to which stages of the child's life are the most important, except within the very early years. Due to data limitations most studies focus on one specific year of investment, and attribute the investment in that given year, and not prior investments, to future development. As investments today build on prior investments, singling out a specific year or a specific period of the child's life does not account for *all* the investments made by working. Instead of examining a specific year, Ruhm (2008) accumulates the years of maternal employment to estimate the effect of the *entire* investment of maternal employment. He does so for two different spells, both starting at the beginning of the child's life, examining the first three and the 10 years (11 for some children) of the child's life on the child's GPA. What is unknown is how maternal employment over a longer time period causally influences child achievement, and we therefore follow the framework by Ruhm (2008).

This paper estimates the causal effect of employment among Danish women on one measure of children's achievement: their academic performance in 9th grade. We take advantage of a unique data set containing detailed household information from over 335.000 Danish children born between 1987 and 1992 who were followed from birth through 9th grade. This allows us to examine the impact of maternal employment over an extended period of time (i.e., a child's first 15 years of life) on an important, long-run outcome: grades in 9th grade. In doing so, we utilize three different methods to address the endogeneity of maternal employment. First, we use the extensive data available from the Danish registers to control for many household and children's characteristics that may affect both children's educational outcomes and maternal employment decisions. Second, we instrument for maternal employment with the local female unemployment rate. Third, we use maternal fixed effects to relate sibling differences in maternal employment experiences to differences in school achievement.

In doing so, we move the literature on maternal employment and child achievement forward. Drawing on a large high-quality data set from a very different social and political context as the Danish with extensive public supports for working parents, we utilize a variety of methods to generate robust associations between maternal employment patterns and children's school achievement in adolescence.

The counterfactual to maternal care differs from country to country, but relatively good substitutes exist in developed countries, and in particular in Scandinavia. The time investment in the child's cognitive development is a complex production function that depends on both the amount of time, and the quality of that time. Despite the mothers' absence, maternal employment benefits the child's cognitive development if the counterfactual care inputs are of a higher quantity and quality than what the mother herself can provide. The mothers' employment provides income so that the family can purchase inputs such as tutoring, books, and computers that may benefit the child's learning. Not only the time the child spends in alternate care, but also the quality of the time she spends with her child outside work will impact the child. She can obtain skills from colleagues that increase the quality of her time with the child, but the opposite may also hold true, that balancing work and family may be stressful and decrease the quality of her investments when she is at home. Empirical findings determine which mechanisms are more prevalent in the given study.

As the U.S, Denmark is characterized by a high degree of female employment, but unlike the U.S., Danish women with children work more than women without children (Statistics Denmark, 2012). Consequently the potential impacts of maternal employment will have large consequences. The potential impacts of maternal employment are most likely driven by the quality of the counterfactual time the child spends away from home. The relative high quality day care in Denmark, along with other support of working families, such as paid sickness absence and paid temporary leave to take care of sick children, serve as a case to find out if a supportive context is enough to offset or even outweigh the impact of maternal time away from her child. Heckman and Cunha (2007) focus mainly on the skill formation on disadvantaged children, but the literature on maternal employment most often find that maternal employment harms children in high SES families more than children in low SES families (Ruhm, 2008; Danzer and Lavy, 2013). This could be due to the poor quality of the counterfactual care compared to a highly educated mother. As for other mechanism that may affect child development, Danish mothers and fathers are among the parents in the world that spend the most time on child care (Bonke, 2009). Both in the U.S and in Denmark, evidence shows that differences in time with children across working and non-working mothers are not large, and that working mothers protect time with their children by cutting back on other areas such as sleep and leisure (Bianchi, 2000; Bonke, 2009). Therefore potential differences are less likely driven by differences in actual direct time investments by the mother, but rather time spent in day care or with the father. Furthermore, the wage distribution is more compressed than in the U.S. why we hypothesize that income is a less important factor as a child-investment in Denmark than in the U.S.

II. Previous Research

This study links maternal, not paternal, employment to child achievement. Evidence suggests that mothers, even those who work outside of the home in the U.S, continue to play the key role of caregivers and managers of their children's time (Bianchi, 2000). Mothers spend more time on child care and housework than do fathers, even in dual-career households (Bianchi, Milkie, Sayer, and Robinson, 2000). Whereas employed mothers perform fewer household and child-related tasks than do those who stay at home, this is not offset by increased time contributions at home from husbands (Cawley and Liu, 2012). As such, there is reason to believe that maternal employment is particularly salient for child development. Indeed, prior research has found that maternal employment was a significant predictor of children's body mass index, whereas that of mothers' partners was not (Morrissey, 2013; Ziol-Guest, Dunifon, and Kalil, 2013). There are exceptions. Greve (2011) finds no overall effect of maternal work hours on child obesity in Denmark, and she attribute the absence of a significant correlation to the high quality Danish day care and paternal contribution in the household.

Only one study (Ruhm, 2008) has examined at least 10 years of childhood, comparing earlier vs. later childhood employment effects on child development, finding that it is employment in the first three years of life that has the strongest, and negative, associations with child cognitive development. Following these finding, this study therefore examine the impact of maternal employment in the first three years, and in the first fiveteen, seperately. Fiveteen years, because we have access to the entire spell from birth to examination in ninth grade at age 15, and cumulative, because the impact of the first years may very well influence the following years of maternal employment and thus child outcomes.

There is no consensus in the literature at which age maternal employment either benefits or harms the child's cognitive development, and at what margin: fulltime, part-time, not employed, or the exact amount of hours worked. Furthermore the different child outcomes are mixed ranging from child health, to cognitive and the wide concept of non-cognitive skills. All measures are intertwined and correlated and used either as a direct outcome or as a mediator for later outcome. Therefore we present the literature on maternal employment (using various definitions) and the effect on several outcomes, all as a measure for the term well-being. We clearly state when the outcome is cognitive skills.

A large body of research in the U.S. examines the consequences of maternal employment for children (Korenman and Kaestner, 2005; Smolensky and Gootman, 2003; Blau and Currie, 2004;

Ruhm, 2000, and Haveman and Wolfe, 1994). Results are mixed, with some studies showing that maternal employment in the first months of life is associated with small, but significant, declines in children's subsequent cognitive outcomes (Hill et al., 2005), particularly when mothers work more than 30 hours per week (Brooks-Gunn et al., 2002). Evidence also suggests that linkages between early maternal employment and lower child health are strongest for children born to more highly educated mothers (Ruhm, 2008).

A child's characteristics may influence both his mother's employment and his own academic performance (Ruhm, 2004). For example, mothers whose children are not doing well may reduce their work hours to invest more time in child-rearing. Indeed, Nes et al. (2013) document linkages between severe behavior problems among children and an increased risk of leaving employment among mothers in Norway. Conversely, mothers with relatively high ability and education may be both more likely to work and to have children who receive good grades in school. Failure to account for unobserved factors that could be linked to both maternal employment and child achievement may bias estimates of the association between maternal employment and children's subsequent academic outcomes.

One way to address the endogeneity of maternal employment is to include a rich set of household characteristics in an ordinary least squares or propensity score regression. Several recent studies taking this approach conclude that there is no association or a negative association between maternal employment and children's cognitive and behavioral development. Four of these studies use National Longitudinal Survey of Youth (NLSY) data from the United States to examine children's verbal and math skills beginning as early as three years of age and continuing thereafter (Baum, 2003; Ruhm, 2004; Berger, Hill, and Waldfogel, 2005; and Ruhm, 2008); Berger et al. (2005) also examine a measure of children's externalizing behavior at age four.[1] Baum (2003) and Berger et al. (2005) focus on whether a woman returns to work within 12 weeks of having a baby in order to match the maternal leave length allowed by the U.S. Family and Medical Leave Act (FMLA), whereas the two studies by Ruhm examine maternal employment during a child's first four (Ruhm, 2004) or 10 (Ruhm, 2008) years of life. A fifth study using this approach finds that the children of women in the United Kingdom who worked full-time in their child's first 18 months of life perform relatively poorly on cognitive tests between the ages of four and seven relative to

[1] The NLSY administers the Peabody Picture Vocabulary Test (PPVT) to children at age three and four, and the Peabody Individual Achievement tests of Mathematics (PIAT-M) and Reading Recognition (PIAT-R) to children at ages five and older.

children of women who worked part-time or not at all (Gregg, Washbrook, Propper, and Burgess, 2005).

Two of the studies mentioned above found that the association between maternal employment and children's cognitive development differed by household characteristics. The effect of maternal employment was more strongly negative among children of highly educated women (Gregg et al., 2005) and women with high socioeconomic status (Ruhm, 2008). Other studies find that children may fare best when mothers work part-time (Brooks-Gunn et al., 2002; Ruhm, 2008; and Hill et al., 2005). This supports the hypothesis of the counterfactual care being the main driver of the negative effects of maternal employment.

One criticism of the above studies is that there may be important unobserved variables, such as the health and cognitive ability of a young child or maternal attributes, which affect a woman's labor supply decisions. Several recent studies address this concern and try to estimate the causal effect of maternal employment on children's cognitive development or academic performance by using a regression discontinuity or differences-in-differences design, maternal fixed effects, or by exploiting exogenous variation in maternal employment.

The studies utilizing econometric methods to address the endogeneity of maternal employment are mixed. Several studies examine linkages between expansions of parental leave policies and child outcomes, looking across countries in Europe and Canada, and considering a wide range of outcomes spanning early childhood into adolescence and adulthood. Several such studies find no linkage to children's cognitive development, health of mother nor child, and later wages (Dustman and Schonberg, 2012; Baker and Milligan, 2010; Rasmussen, 2010; and Dahl, Loken, Mogstad, and Salvanes, 2013). However, Carneiro, Loken, and Salvanes (2011) find that a four-month increase in maternity leave in Norway was associated with a decline in high school dropout rates and an increase in wages, while in Sweden, extended parental leave was associated with improvements in children's educational performance, but only among children with highly educated mothers. One additional month of parental leave entitlement improved test scores by 2 percent of a standard deviation for these children (Liu and Skans, 2010). Additionally, a year-long extension of parental leave in Austria (extending through the child's second year of life) was insignificantly associated with test scores at age 15; however, results differ such that boys with more educated mothers see improvements in age 15 test scores, while boys with less educated mothers demonstrate lower academic outcomes (Danzer and Lavy, 2013).

Finally, using NLSY data and employing maternal fixed effects, Ruhm (2004) finds no statistically significant effect of maternal employment on children's test scores, and James-Burdumy (2005) finds negative, positive, and no effects depending on the test score used and the particular time period when maternal employment is measured (with the strongest negative effects occurring when mothers were employed in the first year of life. One 8 hour additional work day per week in the first year resulted in a 3 percent of a standard deviation decrease in math scores). Bernal (2008) also uses NLSY data, estimating a structural model of children's educational performance and maternal employment, and conclude that full-time maternal employment in the first year reduces children's high school test scores. The majority of studies in this area link maternal employment in the very earliest years of life to child outcomes.

Overall, none of the studies reviewed above conclude that, across all children, maternal employment benefits children's cognitive development. However, evidence suggests that, for disadvantaged children, maternal employment is associated with improved outcomes. For example, Danzer and Lavy (2013) reviewed above, find that children of less educated mothers may benefit from maternal employment. Results from experimental welfare-to-work studies taking place in the early 1990s demonstrate improvements in school performance and reductions in behavior problems among young children when mothers left welfare for work (Gennetian and Miller, 2002; Duncan et al., 2007), yet such programs often had the opposite effect on teenagers (Gennetian et al., 2004). Utilizing a sample of disadvantaged mothers and children, and a variety of methods, Johnson, Kalil, and Dunifon (2012) find that, compared to children whose mothers do not work, those whose mothers work and experience job stability exhibit fewer behavior problems.

Thus, the studies reviewed above suggest that, overall, there may exist small detrimental, or neutral, linkages between maternal employment and child well-being. For disadvantaged children, maternal employment may be associated with positive outcomes.

III. Maternal Employment and Social Policies in Denmark

In this section we discuss the Danish context in which this study takes place. As noted above, the time-related linkages between maternal employment and child cognitive outcomes will be affected by maternal leave policies, the quality of child care available, and work intensity. In several ways, Denmark differs dramatically from the United States along these dimensions.

Maternal employment in Denmark is among the highest in the OECD. The trends in female labor force participation is fairly constant; 74.5 percent in 1985, 78.1 percent in 1990, and 74.1

percent in 1996 (Brewster and Rindfuss, 2000). Among mothers of children between the ages of one and three, labor force participation increased from 70 percent in 2002 to 80 percent in 2008, and from 76 percent to 82 percent for mothers of children between the ages of four and six (Authors' calculations). Additionally, due to collective agreements in Denmark, the standard full-time employment is set to 37 hours per week, however, weekly workhours can very well exceed 37 hours (Greve, 2011). In 1987 Danish women were working 6.340 hours per weekday (31.7 hours per week) on their paid job (Datta Gupta, Bonke and Smith, 2003). On average in 1985, among working women 54.5 percent were working fulltime, increasing to 60.3 percent in 1997. About 25 percent of Danish women worked less than 30 hours per week in 1997 (Auer, 2000; Greve (2011) shows a decrease in Danish women's work hours from 1997 through 2007.

Furthermore, Denmark has a tradition for universal benefits, such as unemployment benefits for ensured and non-ensured, and universal health care. For families with children, Denmark has a generous maternal leave policy. During our sample period (children born between 1987 and 1992), Danish women were entitled to four weeks of paid maternal leave before the delivery; 14 weeks immediately following delivery; and 10 weeks of additional parental leave that could be shared between parents. Following the collective agreements all Danish employees are entitled to paid sickness leave. From 1990 through 2005, the public and private mandatory spending amounted to 0.9-1.2 percent of GDP (OECD, 2008). Danish employees are also entitled to holiday allowance and the right to five weeks of holiday. In 2003 the right increased to six weeks.

Finally, Denmark has a strong system of early care and education (ECE) programs, spending 1.2 percent of its GDP on ECE activities. As a result, Danish households spend on average only eight percent of their income on child care costs; in contrast, the U.S spends 0.4 percent of its GDP on ECE activities, and families spend 19 percent of their income, on average, on child care (Ruhm, 2011).[2] Danish child care is provided at the municipality level, with local governments providing center-based care as well as organizing a system of family-based care in private homes (Datta Gupta and Simonsen, 2010). The staff-to-child ratio in Denmark is among the lowest in the OECD (Datta Gupta, Smith and Verner 2008), indicating a high quality of child care. Between 2001-2006 the child-to-staff ratio in Danish institutions for children younger than 3 years old was fairly constant close to 3.1 children per staff. For preschoolers aged 3-6, in the same time period, the child-to-staff ratio was 5.75 child per staff (Gørtz and Andersen, 2014). As a result of this public

[2] Denmark ranks second out of all OECD countries, behind only Iceland, on percent of GDP devoted to ECE programs.

investment, 63 percent of zero to two-year-olds and 94 percent of three-year-olds are in some type of formal child care in Denmark.

The extensive paid maternity leave time, relatively lower work hours, and generous and high quality ECE system in Denmark suggest that working mothers may face less constrained choices, and therefore the implications of maternal employment for child well-being may differ, compared to those in other countries, such as the U.S. The goal of this paper is to examine the linkages between maternal employment and the academic progress of Danish teens. Because of the large body of work, noted above, suggesting that maternal employment in the first three years of life has particular implications for child well-being, and a lack of studies examining maternal employment across childhood, we perform separate analyses looking at maternal employment in the first three years, and then in the first 15 years, of childhood. Due to the evidence, cited above, of non-linear associations between maternal work hours and child well-being, we examine how child GPA differs based on varying intensity of maternal work. Finally, because of studies indicating that any detrimental linkages between maternal employment and child well-being are concentrated among more advantaged mothers (Ruhm, 2008), we perform some analyses stratifying our sample by maternal education.

IV. Data

Since 1980, Statistics Denmark has collected and recorded data on all individuals living in Denmark. The registers contain all public transfer payments between individuals and the federal/municipal government, such as income taxes, parental leave benefits, day care subsidies, unemployment benefits, and pension payments. Danish law requires individuals to inform the federal government of their residential location; individuals who fail to comply are not allowed to have a bank account or receive benefits from the state. As a result, the registers contain information on the entire Danish population, with attrition only due to deaths or migration.

Our sample includes all children born in Denmark between 1987 and 1992. We focus on this time period because the registers contain grade point averages (GPA) in 9th grade for these children, which for most Danish children is when they are 15 years of age. We use three methods to examine the associations between maternal employment and child achievement: the ordinary least squares, instrumental variables, and maternal fixed effects models (the latter of which, as described below, are limited to mothers with at least two children).

Measures

Children's Grades in School. The dependent variable in our analysis is a child's GPA in 9th grade. For the GPA measure to be comparable between students, we use a GPA from four 9th grade courses that all Danish students take: Math, English, Science, and Danish. We include all the grades in each subject, and condition on at least one grade in each subject for a child to be included in the the sample. Although some Danish schools offer additional subjects such as German, French and Chinese, we exclude these grades from the GPA measure because these subjects are taught for only a few hours per week and only in later grades. Each school reports the grades annually to the Ministry of Education, which in turn forwards the information to Statistics Denmark.

The Danish grading system is different from the American system, with scores ranging from a low of zero to a high of 13. Zero through five are considered to be failing grades. During the period we examine, the scoring system changed. We use guidelines from the Ministry of Education to transform the old scores to be consistent with the new ones.[3] To make the estimates comparable between countries and studies, we standardize the GPA with a mean of zero and a standard deviation of 1.

Maternal Employment. Our key independent variables measure maternal employment. For the period we examine, the Danish registers do not contain actual annual working hours. However, each tax payer must pay a certain mandatory pension-saving (ATP) which is based on thresholds of working hours.[4] The worker pays ATP according to the thresholds: working 10-19 hours per week, 20-29 hours per week or more than 29 hours per week. If she works less than 10 hours per week, she does not pay ATP. From these pension payments we deduce a mother's annual work hours. We create four categorical measures of work hours: working 0-9 hours per week (including those who are not working; this category is also referred to as "unemployed" and serves as the omitted category in our analysis); between 10 and 19 hours per week; between 20 and 29 hours per week; or 30 hours or more per week. If, for example, a person is working 15 hours per week on average in a given year, the person will pay taxes according to the second category – working between 10 and 19 hours per week. Workers only pay the tax if they are employed more than half a calendar year; therefore people who worked but did so for less than half of the year are classified as unemployed (working fewer than 10 hours per week). We assume that people working less than 10 hours per

[3] The transformation is: new -3=0 old, new 0=3 old, new 2=6 old, new 4=7 old, new 7=8.5 old, new 10=10 old, new 12=11 old (Guidelines from Ministry of Education).

[4] It is voluntary for self-employed to pay ATP. If they pay ATP, they will be part of our sample. If not, and if they do not receive any unemployment benefits, they are not in the sample. In 1995, the self-employment rate for women, including non-mothers, was 3.8 (Carrasco and Ejrnæs, 2003) and 3.7 in 2009 (Statistics Denmark, 2011).

week on average annually are unemployed. Additional analyses using supplemental survey data, not shown here, indicate that the vast majority of mothers classified as working 0-9 hours per week are indeed not working.

Time spent on maternity leave is measured in the registers in the same way as prior employment or unemployment; that is, if a mother is employed full-time prior to a delivery, she will also pay ATP during her maternity leave. The opposite is true if a woman was unemployed before a delivery. Because we measure maternal employment during a full calendar year, this issue comes up in many surveys that collect data at a certain point in time. During our sample period (children born between 1987 and 1992), Danish women were entitled to less than one-half a year of maternity leave: four weeks before the delivery; 14 weeks immediately following the delivery; and 10 weeks of additional parental leave that could be shared between spouses. As a consequence, the estimated working hours of mothers who used all of their maternity leave during this time period will be an average of her pre- and post-birth working hours over the course of the entire calendar year. If she is on maternity leave the first half of the year (but worked before that) and becomes unemployed for the second half, the ATP payments we observe will be based on her employment for six months and her unemployment for six months. Although measurement error is present, we believe it will be small, and only within the given calendar year. As for women who spend three months on leave in one calendar year, and three months in the next, they will be categorized according to the category they belong to for the remaining nine months of those years. Of course, maternity leave is staying at home, not working. But because the maternity leave period does not exceed six months, we can justify that the observed work hours for a given year will be based on the largest part of the year.

We are interested in estimating the effect of a woman's work effort during her child's first three years and first 15 years of life on the child's grade point average at age 15. The work hours follow the calendar year.[5]To estimate a woman's average weekly hours worked during a child's first 15 years of life (and similarly for the first three years of life), we create four indicator variables for each mother for each calendar year: 1) unemployed; 2) working 10-19 hours per week; 3) working 20-29 hours per week; and 4) working 30 or more hours per week.

[5] To examine the differences in treatment for children born early and late in the year, we stratify the sample on birth month in the *supplementary analyses* section.

In some regression specifications we include the percentage of her child's first 15 (or three) years a woman spent in each of these work-hour groups, as follows:

$$av\_work\_0 = \frac{\sum years\ unemployed}{15}$$

$$av\_work\_1 = \frac{\sum years\ working\ 10 - 19\ hours}{15}$$

$$av\_work\_2 = \frac{\sum years\ working\ 20 - 29\ hours}{15}$$

$$av\_work\_3 = \frac{\sum years\ working\ 30 + hours}{15}$$

We observe the full 15-year employment history for 56 percent of the mothers. About 15 percent of the women have missing work hours information in a given year, and this missing rate is fairly constant over the 1987 to 2007 time period when the children in the sample are less than 16 years old. People who did not make any tax payments are included as being unemployed in that given year. Missing employment data could be due to many reasons. This could occur for adults out of the labor force due to illness, early retirement, lack of eligibility for benefits or volunteer unemployment. It could also be immigrants who have not yet formalized their citizenship or families working abroad for a period of time.

In some regression specifications we include continuous measure of a woman's average hours worked per week over the first 15 years (or three years) of her child's life. A continuous work hour variable is necessary for the two-stage least squares specification (described below). Because the work hour categories are intervals, for purposes of estimating a continuous employment variable we assign hours in the middle of an interval (e.g., 15 hours per week for women working between 10 and 19 hours per week in a particular year). We assume that women working 30 or more hours per week actually worked 35 hours, and women working less than 10 hours per week were actually unemployed.

Figure 1 depicts the pattern of maternal employment by a child's age for the first 15 years of a child's life. About 65 percent of Danish mothers worked when their child was one year old, with the majority of the employed mothers working more than 30 hours per week. The percentage of mothers working more than 30 hours per week increases as children age, from 41 percent when a child is one, to 50 percent when a child is five, to 61 percent when a child is 10, and to 69 percent when a child is 15 years old. The percentage of non-working mothers decreases as children grow up; by age 15, 20 percent of the mothers are unemployed or working less than 10 hours per week. Ten percent of the mothers work part-time (10 to 19 hours or 20 to 29 hours per week) when children are 15 years old, in contrast to 22 percent when children are one.

Controls. In all regressions we include birth-year and birth-month fixed effects to control for national GPA trends in Denmark and age effects within a grade. In most regressions we include a more complete set of control variables, including: the child's gender and ethnicity, the mother's age when the child was born, the mother's education in the year before the child was born, the percentage of the child's first 15 (or three) years spent living in a nuclear family, and household income in the year before the child's birth. We measure a woman's education with six separate indicator variables: less than a high school degree, high school degree only, vocational school, short-term further education, bachelor's degree, and master's degree or more. We use mothers' age when the child was born, her education, and household pre-birth income as proxies for pre-birth ability. We estimate a household's net income based on the parents' wages and other cash transfers, measured the year prior to a child's birth as an additional control for the parents' ability. Income is measured in thousands of Danish Kroner, where a thousand Kroner is about $170. In unreported regressions we include the average income a household received over their child's first three (or 15) years rather than their income prior to the child's birth.

The presence of a father in the household could affect both the academic achievement of the child and how much the mother chooses to work. Therefore, we control for the percentage of the child's first 15 (or three) years the child lived in a nuclear family. We consider a child to be living in a nuclear family when the child and the biological parents live at the same address. Although the registers do not record a person's race, they do indicate whether a person immigrated to Denmark or whether the previous generation of the family immigrated to Denmark. We include indicators for whether a child is a first- or second-generation immigrant.

Angrist, Lavy and Schlosser (2010) find a large negative correlation between family size and child schooling in Israel. However, instrumenting for family size using twin births as an instrument, these correlations disappear. Black, Devereux and Salvanes (2005) come to the same conclusion in Norway. Examining other channels in family structures on educational attainment, they moreover find that birth order is more important for the child's educational outcome than family size. Having siblings might reduce maternal working hours, but it does not necessarily give more maternal care with the child, and this time investment could be different for each number in the birth order. Therefore we also include an indicator for a child's birth order in his/her family.

A woman's decision regarding whether and how much to work following a child's birth may be affected by the child's health (Powers, 2001). Therefore, from the fertility register records we create indicator variables for children who were born prematurely (born earlier than the 37^{th} week), who had low birth weight (less than 2,500 grams), or who have chronic health conditions. The registers include the World Health Organization's (WHO) International Classification of Diseases (ICD) codes for all hospitalizations in Denmark. To classify a child as having a chronic health condition, we use the same definition of chronic diseases as Christoffersen et al. (2003), including diseases that are non-psychological and persistent. Because the codes do not specify disease severity, we classify children as being chronically ill if they were hospitalized for one or more of the designated chronic diseases between 1987 and 1996, a period when all of the children in our sample were at least five years old.

Finally, our full set of control variables also includes a set of variables recording the proportion of a child's first three years or first 15 years that he/she spent in each municipality. If women (and their children) never moved, this would be equivalent to including a full set of municipality fixed effects. Including the municipality variables allows a child's performance in school to be affected by municipality-specific factors, such as educational quality, and allows this effect to be proportional to the amount of time a child spent in each municipality. Once we omit households with missing GPA values[6], 335,199 children remain in the analytic data set.

In Table 1 we examine employment patterns for women in the sample. The table reports the percentage of women by the number of years (out of 15) they spend in a particular work-hour category. Sixteen percent of the mothers worked 30 or more hours every year when their child was under the age of 16, and almost one half worked 30 or more hours in at least 10 of these 15 years. Thirty-five percent of the women worked fewer than 10 hours (including zero) in at least one of the 15 years, although only one percent worked fewer than 10 hours in 10 or more of the 15 years. Many women worked 10-19 or 20-29 hours at some point when their child was under the age of 16.

We report sample statistics in Table 2, separately by the average number of hours per week that a mother worked during her child's first three years of life. The children of women who work relatively long hours have relatively high GPAs. Women who work relatively long hours are also more educated, are more likely to live in Copenhagen, tend to have children at an older age, have children who spend more time in their nuclear family, and tend to have healthier children. These

[6] Another interesting outcome is passing the GPA, including the children that did not hand in any tests. This is however beyond the scope of this paper.

differences emphasize the importance of controlling for both children's and mother's characteristics in the regression models, and of trying to control for differences across employment categories in unobserved characteristics that affect children's GPAs. We elaborate on our approach to this challenge in the following section.

V. Method

Our basic empirical approach is to regress a child's GPA at age 15 on his/her mother's employment status and work hours during the first three years of the child's life, and separately during the first 15 years of the child's life. The regression coefficient on the maternal employment variable(s) would indicate whether maternal employment is positively or negatively associated with children's school outcomes. In the ordinary least squares regression models, we regress child c's GPA in year t on the mother m's employment (H) during the child's first three or 15 years of life, and various control variables (X) for the child and mother:

$$(1)\quad Y_{c,m,t} = \beta_0 + \beta_1 H_{m,15} + \beta_2 X_c + \beta_3 X_m + \varepsilon_{c,m,t}$$

We report results from analyses in which H represents a set of variables measuring the percentage of the three or 15 years a mother spent working 0-9 hours a week (omitted category), between 10 and 19 hours, between 20 and 29 hours, and 30 or more hours. We also report a similar set of results using a continuous measure of the mother's work hours in her child's first three or 15 years of life. Standard errors are clustered by household to account for the fact that our data contain siblings.

In some unreported regressions, we also include a control for household income over the child's first 15 years (and three years) to test whether a mechanism through which maternal employment may influence children is through the additional financial resources such employment brings.

As noted above, a mother's ability and education may be associated both with her work hours, and with her child's grades in school. On the one hand, mothers with greater ability and education may achieve higher status jobs with longer work hours, and also have children with positive educational outcomes (due to differences in genetics, parenting, or other factors between mothers of different abilities and education). On the other hand, mothers with lower ability and education may need to work longer hours in order to make ends meet, and may also have children doing poorly in school. Without controlling for these potentially unobserved maternal characteristics, β_1 is likely to be biased. The sample statistics in Table 2 confirm that in our data set this challenge is real. Children whose mother averaged more than 30 hours of working per week when the child was under the age of 15 have a 9^{th} grade GPA that is 0.6 points higher than children whose mother averaged fewer than 10 hours of work per week.

We address the endogeneity of maternal employment in three different ways. First, we use the extensive data available from the Danish registers to control for many household and children's characteristics that may affect both children's educational outcomes and maternal employment decisions. To demonstrate the importance of controlling for a rich set of characteristics, we first present regression results that control only for birth year and birth month. In subsequent regressions we add the full set of control variables.

Our second method of addressing the endogeneity of maternal employment is to instrument for a mother's hours worked in year t with the previous year unemployment rate in her municipality among women with the same education level. This approach is similar to that of Greve (2011), who uses the Danish register data to examine the effect of maternal employment on child obesity. She shows that a woman's working hours in 1999 (when the children in her sample are three years old) is negatively correlated with the unemployment rate in her municipality (controlling for income and education). Before a municipal merger reform in 2007 reducing the number of Danish municipalities to 98, Denmark consisted of 271 municipalities, and the unemployment rate ranged from 2.2 to 11.0 percent across municipalities in 1999. When examining employment over 15 years of her child's life, we first calculate the average female unemployment rate in a woman's municipality, stratified by education level, over the 15-year period when her child was under the age of 16. In the first stage of a two-stage-least-squares regression, we regress a woman's average hours worked per week over the 15-year period on the average unemployment rate in her municipality (for women with the same education level) for the same time period, along with the other controls, including the municipality fixed effects noted above.[7] When examining maternal employment in the first 3 years of the child's life, we use a similar method, focusing on the unemployment rate in her municipality over the first three years.

Identification in the IV model comes from variation in hours worked due to variation in the gender- and education-specific unemployment rate within a municipality over time (because municipality fixed effects are included), and thus to variation in employment opportunities and wages. As we show below, the instrument is negatively correlated with the endogenous variable, maternal work hours, and highly significant statistically.

[7] For women who live in different municipalities over this time period, we use the relevant unemployment rate for each municipality-year.

The key identifying assumption is that, after controlling for our rich set of child and household characteristics, variation in the female unemployment rate within municipalities over time does not directly affect children's subsequent GPAs other than through the effect on female labor supply. Bernal (2008) makes a similar assumption. One concern is that municipalities with high unemployment rates may have low-quality schools, and thus low student GPAs, due to reduced tax revenues and reduced education funding or due to unobserved characteristics of parents and children who live in such locations. This is not likely to be as great a concern in Denmark because, although local tax revenues are an important financing source for education, at a national level, Denmark provides large federal transfers from rich to poor localities. Additionally, our models include the municipality fixed effects, which are a full set of municipality exposure variables to account for time-invariant municipality-specific factors that affect student outcomes, such as the inherent quality of a municipality's school system. A limitation of any IV method is that estimates may only be relevant for women who are "compliers"—i.e., those whose labor supply is affected by the local unemployment rate. As discussed in more detail below, a very small proportion of our sample consists of "compliers", limiting our ability to generalize from our IV results.

Our third method of addressing the endogeneity of maternal employment is to use maternal fixed effects to control for differences between households in their children's abilities. The maternal employment coefficient in this specification is identified by differences between siblings in a mother's work intensity (recall that women in the fixed effects analyses only include mothers who gave birth to two or more children between 1987 and 1992). There are a variety of reasons why a mother's work intensity may vary across siblings; it is possible that these same reasons are also correlated with children's GPA in adolescence. For example, mothers may reduce their work hours for a child with health problems (Neidell, 2000; Ruhm, 2004). Additionally, other unobservable household or individual changes may be associated both with variation across siblings in mothers' work hours and in sibling differences in educational achievement. While we include controls for a child's health status, described above, to the extent that changes in unobserved individual or household-level variables are important, however, we would expect the maternal employment coefficients from this method to be biased. Relative to the IV method, therefore, the maternal fixed effects models provide less robust evidence of a causal effect of maternal work hours on a child's subsequent educational performance.

Another potential concern is that the employment decisions for women who have two or more children in a five-year period may differ from those women who have only one child, such that the results from the maternal fixed effects model are not generalizable. In Figure 1 we display the 15-year pattern of employment for all women in our sample period (at the top with missings included as unemployed and at the bottom with missings in a separate category).

VI. Results

We now turn to results of the analyses linking maternal work hours to a child's GPA at age 15. In Table 3 we present results of analyses using ordinary least squares and maternal fixed effects regressions to examine these associations, focusing on maternal employment in the first three years of a child's life. First we report OLS regression results that include a set of variables measuring the percentage of the child's first three years that a woman spent working between 10 and 19, 20 to 29 hours, and 30 or more hours per week relative to the omitted category (percentage of years working fewer than 10 hours per week, including not working at all). This specification allows for a non-linear effect of work hours on children's outcomes. When we only include birth-year and birth-month fixed effects in column one, the coefficients measuring the percentage of the child's first three years that a woman spent working are all positive and significant relative to the omitted category. The coefficients on the three employment variables increase monotonically as average working hours increase.

In the second column of Table 3 we include the full set of control variables previously described. The inclusion of the control variables in column two has substantially decreased the employment coefficients. The results indicate that working 10-19 hours per week while a child is under four years of age is associated with a 0.135 standard deviation increase in child GPA at age 15, compared to working less than 10 hours per week. The coefficients on the percentage of time working 20-29 hours and 30 or more hours per week have a smaller magnitude between 0.083-0.094 standard deviation. In Appendix 1 we report coefficient estimates for the full set of control variables that are not reported in Table 3.

In column three we include a continuous hours-worked variable in order to facilitate comparisons to the IV specifications discussed below. The coefficient on a mother's average hours worked per week is positive, significantly different from zero, and indicates that an additional hour worked per week between ages zero and three is expected to increase average GPA by 0.29 percent of one standard deviation. The predicted magnitude of working 35 hours per week on a child's GPA is 0.105, which is very similar to the final coefficient reported in column two (0.094).

In the final two columns of Table 3 we report coefficient estimates from a maternal fixed effects regression. In these specifications we also include the full set of control variables. The coefficient on the linear hours-worked variable in column 4 is not significantly different from zero. When we include categorical employment indicator variables in column five, only the coefficient on 10-19 hours is positive and significant, and the magnitude is much smaller than found in the OLS regressions (column 2).

In Table 4 we present results from a similar set of regressions examining a mother's employment history over the child's first 15 years of life. The coefficients on maternal work hours remain positive and significant in the OLS regressions, and the magnitudes are larger than when examining maternal employment only in the first three years, suggesting that the effect of maternal employment accumulates for many years over a child's life, rather than being focused only on a child's first few years.

In column two of Table 4, which includes a full set of control variables, the results show that the percentage of time a mother worked 10-19, 20-29, and 30+ hours per week over the child's lifetime is associated with a 0.041, 0.302 and a 0.225 standard deviations increase from the average GPA at age 15, respectively, compared to working less than 10 hours per week. A child of a woman who worked 30 or more hours per week is predicted to have a five times higher percentage of one standard deviation from the mean GPA relative to the child of a woman who worked fewer than 10 hours per week. As in the first set of regressions, the maternal employment coefficients are considerably smaller when we include detailed characteristics of the child and mother. In column three the coefficient on the linear hours-worked variable is positive and significant.

In the final two columns of Table 4 we report coefficient estimates from maternal fixed effects regressions. As for the child's first three years, only the coefficient on 10-19 hours is positive and significant.

We next present results in Table 5 from the two-stage least squares models in which we instrument for a woman's work hours with the female, education-specific unemployment rate in her municipality during the first three or 15 years of her child's life.[9] Column one and three show the reduced form estimate, regressing female, education-specific unemployment rate on the standardized GPA. The results show a significant negative correlation between the instrument and the GPA, indicating that higher female, education-specific unemployment rate decreases the GPA.

[9] Because we lag the unemployment rate one year, we actually use the year prior to the child's birth plus the first two years, and so forth for the 15-year analysis.

For the three year measure, a 1 percentage point increase in the unemployment rate decreases the GPA with 0.06 percent of one standard deviation, and for the 15 year measure, 0.2 percent of one standard deviation. We argue that the instrument works through maternal employment and not directly, but we cannot statistically test this exclusion restriction. The main concern is that high local unemployment also correlates with poor maternal childrearing ability and poor labor market attachment. If this is the case, the instrument is directly affecting the child's GPA, and the results will be biased.

To further test the exclusion criterion, we use survey data to estimate correlations between our instrument and one aspect of maternal child rearing abilities: maternal mental health. In Appendix 2 we estimate the correlation between maternal health and unemployment rates, stratified by education, in above and below median unemployment municipalities. The estimates are small and insignificant. From the positive reduced form estimate and the results in appendix 2, we argue that the exclusion restriction is met, or put differently; we cannot find evidence to suggest that the exclusion criterion is *not* met.

The local unemployment rate is assumed to affect a child's GPA solely through its impact on a woman's labor supply decisions. In each of the two first-stage regressions (shown at the bottom of Table 5, columns two and four), the coefficient on the local unemployment rate is negative as expected and significant, and the F-statistics are very large[10]. A one-percentage point increase in the female, education-specific unemployment rate in a woman's locality is associated with a decrease of about 0.32 hours worked per week.

Looking at the top panel of Table 5, the second-stage maternal employment coefficients are positive, significant, and larger than the coefficients from the comparable OLS specifications reported in Tables 3 and 4. In the second column of Table 5, each additional hour per week of maternal employment during the first three years of a child's life is associated with 1,6 percent of one standard deviation increase from the average 9th grade GPA. The child of a woman who worked an average of 30 hours per week, while the child was under the age of four is predicted to have a GPA that is 0.48 standard deviation higher than an otherwise similar child whose mother did not work at all.

[10] The F-statistics are very large, and we believe it is a consequence of a full-population study. In the subsample analysis, the F-statistics are reduced by about 50 percent. The standard errors generally become smaller as the sample size increase. Running an estimation including only work hours and the instrument, the F-statistics change to close to zero, namely 0.10.

In the final two columns of Table 5 we report the IV results when analyzing women's employment histories over their children's first 15 years of life. As with the OLS results, the maternal employment coefficients are about three times larger in the 15-year versus the three-year time period. Based on the specification in column four of Table 5, the child of a woman who worked an average of 30 hours per week while the child was under the age of 16 would be predicted to have a GPA 1.53 standard deviations higher than an otherwise similar child whose mother did not work at all.

It is important to note that the IV model estimates a Local Average Treatment Effect (LATE), capturing the effect of maternal employment only for children of the "compliers"—women whose labor supply decisions were affected by the female- and education-specific unemployment rate in their localities (Angrist and Pischke, 2009). The IV model is not able to estimate effects for women whose labor supply decisions were not affected by the local unemployment rate. Supplementary analyses (Appendix 3) shed light on the population of compliers in our IV analysis. Results indicate that compliers (i.e., women whose labor supply is affected by the local unemployment rate) are less educated and have lower incomes compared to the sample average. Overall, compliers are a very small proportion of working women, with only 4.2 percent of less educated women induced to work by local labor market conditions, and only 1.3 percent of more educated women induced to do so. Thus, our IV estimates are identified off of a small sub-sample of women who are compliers. Angrist and Pischke (2009) points out that the instruments used by Acemoglu and Angrist (2000) to instrument for high school graduation, of students graduating less than 2 percent did so because of compulsory attendance laws. The same is the case for the local demand for female labor in this study.

Supplemental Analyses

As noted above, a key mechanism through which maternal employment might predict children's well-being is via the increased income such employment brings. In analyses not shown, we added a measure of contemporaneous household income measured at the same time as maternal employment (the first 3 or first 15 years) to our regression models. Doing so did not change the results presented here (results available upon request).

Additionally, previous research indicates that children of less advantaged women may derive more benefits from maternal employment than other children. Table 6 showing 2SLS results for analyses for different subsamples, we stratified the sample by education, performing IV analyses for women with less than a high school degree (low education) and those with a high school degree or more (high education). The maternal employment coefficients are strictly higher and significant for low-education mothers than high-education mothers.

Gender of the child seems to play a large role within education outcomes generally, and boys could potentially be more affected by maternal employment than girls. However, stratifying on gender of the child, the IV results do only show small differences between girls and boys indicating that boys experience slightly higher GPA's than girls, if their mothers work longer hours.

As the analyses build on yearly work hours, children born in January might experience a different treatment than children born in December. To examine this further, we stratify on children born January through June and July through December. As we would expect and the results show in Table 6, children born later in the year benefit more from a full-time working mother than children born earlier.

A robustness test of the results, stratifying on different subjects: Math, Danish, English and Science the results show that the estimates of the overall results are robust. As Table 6 shows, for all four subjects, the IV-results show positive, significant results, ranging from the lowest estimate for Science (0.010 with standard error 0.002) to the highest estimate in English (0.021 with standard error 0.003).

VII. Conclusions

In this paper we estimate the effect of maternal employment on children's educational outcomes using three different methods. Using detailed data from the Danish administrative registers on 335,199 children born between 1987 and 1992, we examine the association between mothers' work hours during the first three years of the child's life, and separately during the first 15 years of the child's life, on children's grades in ninth grade. In the OLS and IV models, maternal employment is associated with a significant increase in children's grades. Results from our IV model indicate that the child of a woman who worked 30 hours per week while her child was under the age of four is predicted to have a GPA that is 0.48 standard deviation higher than an otherwise similar child whose mother did not work at all. These results stand in sharp contrast to much of the previous research in this area, which generally finds no relationship or a negative relationship between maternal employment and a child's cognitive development or academic performance.

The effects of maternal employment are larger when we examine a mother's employment history over the child's first 15 years of life, indicating that the maternal employment effect is cumulative over many years, rather than being focused only on a child's first few years. This stands in contrast to other work in the U.S. context, which has found that employment in the earliest years of life are more strongly (and negatively) associated with child cognitive outcomes than later employment (Ruhm, 2008; Baum, 2003; Brooks-Gunn et al., 2002; Hill et al., 2005). Our results suggest that it is the accumulation of maternal employment experiences over childhood that matters most (and in a positive direction) for children's success in school.

The analysis also indicates that the positive coefficient on maternal employment does not appear to be driven by increased household earnings. Including a household's contemporaneous income does not materially change the coefficient estimates. This result is consistent with many U.S. studies and suggests that rather than being attributable to the additional resources mothers' bring into the household, maternal employment influences children's cognitive development through other channels. Although our data set is not well-suited to test the possible mechanisms, possible explanations include increased access to social support among employed mothers, access to high-quality child care, and relatively cheap available leisure activities for young people in Denmark. Future work is needed to better understand the mechanisms that may link maternal employment to improved child outcomes.

The magnitude of the effect of maternal employment differs across the three models, with the maternal fixed effects model producing the smallest (and oftentimes insignificant) estimates and the IV model the largest. While we are not able to fully disentangle why our estimates differ across models, as noted above, we have less confidence in the maternal fixed effects as they are most likely to be biased by unobservable factors that may differ between siblings. A key point, however, is that in none of the three models presented here do we find evidence of a significant negative association between maternal employment and children's grades.

In identifying a robust positive association between maternal employment and children's achievement, this study stands in contrast to the larger body of work in this area. One possible explanation for our divergent findings is the Danish context in which this study takes place. Denmark differs from the U.S. in a variety of ways. Most Danish mothers work and do so full-time (37 hours per week in Denmark; Greve, 2011). Such employment is facilitated by generous policies such as high quality child care and paid maternity leave. As noted by Greve (2011), this suggests that, perhaps in other contexts, child well-being suffers not because of maternal employment per se, but rather the lack of public support for maternal employment. By examining the consequences of maternal employment in a vastly different context than the U.S., the current study allows for a more nuanced understanding of how and why it might be associated with children's development. Of the previous econometric work in this area looking at child educational outcomes, only one (Rasmussen, 2010) used Danish data; that study found no linkages with long-term educational outcome.

This paper has some limitations that should be noted. As discussed above, we are not able to shed light on the channels through which maternal employment improves children's school outcomes. While the Danish register data has several advantages, particularly its large sample size, it does not contain information on family processes and other factors that may account for the associations we observe. Additionally, we are forced to infer women's work hours from their pension payments. As noted above, comparisons between our assumptions and direct assessments of women's work hours from survey data increase our confidence in these inferences; however, a possibility remains that, for some women, work hours are mis-measured. Finally, while we employ a variety of methods to address our research question, no method is without its limitations. For example, our IV results may not generalize beyond the small sample of women whose employment is responsive to local labor market conditions, and our maternal fixed effect results are estimated from women who experienced changes in their labor supply across children. Importantly, however, in none of our models was there evidence that maternal employment is linked to detrimental outcomes for children. Instead, looking across all of our models, evidence suggests positive associations between maternal employment and the academic progress of children in Denmark.

Taken together, then, this paper presents evidence of a positive linkage between maternal work hours and the educational attainment of Danish teens. These associations are robust across a variety of estimation methods and not accounted for by household income. These findings diverge from the larger body of work linking maternal employment to child outcomes, most of which finds neutral or small negative associations between maternal work hours and child cognitive development. More work is needed to better understand how and why the influence of maternal employment on children may differ in cross-national contexts. However, the current study provides strong evidence, across multiple models, that, for Danish women, maternal employment is associated with improved outcomes for children.

References

Acemoglu, Daron, & Joshua Angrist, 2000. How Large are the Social Returns to Education? Evidence from Compulsory Schooling laws. *National Bureau of Economics Macroeconomics Annual 2000*, 9-58. MIT Press, Camibridge, Mass.

Angrist, Joshua D., & Jorn-Steffen Pischke, 2009. Instrumental variables in action: sometimes you get what you need. Mostly Harmless Econometrics: An Empiricist's Companion, 113-220.

Angrist, J., Lavy, V., & Schlosser, A. 2010. Multiple experiments for the causal link between the quantity and quality of children. *Journal of Labor Economics*, *28*(4), 773-824.

Auer, Peter.2000. Employment revival in Europe: labour market success in Austria, Denmark, Ireland and the Netherlands. International Labour Organization.

Baker, Michael, & Kevin Milligan. 2010. "Evidence from Maternity Leave Expansions of the Impact of Maternal Care on Early Child Development." *Journal of Human Resources* 45(1):1-32.

Baum, Charles L. 2003. "Does Early Maternal Employment Harm Child Development? An Analysis of the Potential Benefits of Leave Taking." *Journal of Labor Economics* 21(2):409-448.

Berger, Lawrence M., Jennifer Hill, & Jane Waldfogel. 2005. "Maternity Leave, Early Maternal Employment and Child Health and Development in the US." *The Economic Journal* 115:F29-F47.

Bernal, Raquel. 2008. "The Effect of Maternal Employment and Child Care on Children's Cognitive Development." *International Economic Review* 49(4):1173-1209.

Bianchi, Suzanne. 2000. "Maternal Employment and Time with Children: Dramatic Change or Surprising Continuity?" *Demography 37*:401-414.

Bianchi, Suzanne, Melissa Milkie, Liana Sayer, & John Robinson. 2000. "Is Anyone Doing The House-Work? Trends in The Gender Division of Household Labor." *Social Forces* 79:191–228.

Black, S.E., Devereux, P.J. and Salvanes, K.G., 2005. The more the merrier? The effect of family size and birth order on children's education. *The Quarterly Journal of Economics*, pp.669-700.

Blau, David, & Janet Currie. 2004. "Preschool, Day Care, and Afterschool Care: Who's Minding the Kids." National Bureau of Economic Research Working Paper #10670.

Bonke, J., 2009. "How much time and money parents spend on their children" (in Danish: Forældres brug af tid og penge på deres børn). Rockwool Foundation Research Unit and University Press of Southern Denmark.

Brewster, Karin L., and Ronald R. Rindfuss. 2000. "Fertility and women's employment in industrialized nations." *Annual review of sociology*: 271-296.

Brooks-Gunn, Jeanne, Wen-Jui Han, & Jane Waldfogel. 2002. "The Effects of Early Maternal Employment on Child Cognitive Development." *Demography* 39(2):369–392.

Bureau of Labor Statistics, 2014. Women in the Labor Force: A Data Book. Available online at: http://www.bls.gov/opub/reports/cps/womenlaborforce_2013.pdf

Carneiro, Pedro, Katrine Løken, & Kjell G. Salvanes. 2011. "A Flying Start? Maternity Leave Benefits and Long Run Outcomes of Children." IZA Discussion Paper #5793.

Cawley, John., & Feng Liu. (2012). Maternal employment and childhood obesity: A search for mechanisms in time use data. *Economics & Human Biology*, *10*(4), 352-364.

Christoffersen, Mogens Nygaard, Henrik Day Poulsen, & Anne Nielsen. 2003. "Attempted suicide among young people: Risk Factors in a Prospective Register Based Study of Danish Children Born in 1966." *Acta Psychiatrica Scandinavica* 108:350-358.

Cunha, Flavio and James Heckman. 2007. "The Technology of Skill Formation." *American Economic Review*, 97(2): 31-47.
Dahl, Gordon B., Katrine Loken, Magne Mogstad & Kari Vea Salvanes. 2013. "What is the case for paid maternity leave?" NBER Working Paper 19595.
Danzer, Natalia & Victor Lavy. 2013. "Parental leave and children's schooling outcomes: Quasi-experimental evidence from a large parental leave reform". NBER Working Paper 19452.
Datta Gupta, Nabanita., Jens Bonke, and Nina Smith. 2003. Timing and flexibility of housework and men and women's wages. *Institute for the Study of Labor (IZA), IZA Discussion Papers* 860.
Datta Gupta, Nabanita & Marianne Simonsen. 2010. Non-cognitive child outcomes and universal high quality child care. *Journal of Public Economics* 94, 30–43.
Datta Gupta, Nabanita, Nina Smith, & Mette Verner. 2008. "The Impact of Nordic Countries' Family Friendly Policies on Employment, Wages, and Children." *Review of Economics of the Household* 6:65-89.
Duncan, Greg, & Jeanne Brooks-Gunn. 1997. Consequences of Growing Up Poor. New York: Russell Sage Foundation.
Duncan, G. J., Dowsett, C. J., Claessens, A., Magnuson, K., Huston, A. C., Klebanov, P., & Japel, C. (2007). School readiness and later achievement. *Developmental psychology*, *43*(6), 1428.
Duncan, Greg J., K.M. Ziol-Guest, & Ariel Kalil. 2010. "Early Childhood Poverty and Adult Attainment, Behavior, and Health." *Child Development* 81:306-325.
Dustmann, Christian, & Uta Schonberg. 2012. "Expansions in Maternity Leave Coverage and Children's Long-Term Outcomes." *American Economic Journal: Applied Economics* 4(3):190-224.
Gennetian, Lisa A., Greg Duncan, Virginia Knox, Wanda Vargas, Elizabeth Clark-Kauffman, and Andrew S. London. 2004. How welfare policies affect adolescents' school outcomes: A synthesis of evidence from experimental studies. *Journal of Research on Adolescence,* 14: 399–424.
Gennetian, Lisa A., and Cynthia Miller. 2002. Children and welfare reform: A view from an experimental welfare program in Minnesota. *Child Development*, *73*(2), 601-620.
Gørtz, Mette, and Elvira Andersson. 2014. "Child-To-Teacher Ratio And Day Care Teacher Sickness Absenteeism." *Health economics* 23(12): 1430-1442.
Gregg, Paul, Elizabeth Washbrook, Carol Propper, & Simon Burgess. 2005. "The Effects of a Mother's Return to Work Decision on Child Development in the UK." *The Economic Journal* 115:F48-F80.
Greve, Jane. 2011. "New Results on the Effect of Maternal Work Hours on Children's Overweight Status: Does the Quality of Child Care Matter?" *Labour Economics* 18:579-590.
Haveman, Robert, & Bobbi Wolfe. 1994. "The Determinants of Children's Attainments: a Review of Methods and Findings." *Journal of Economic Literature* 33:1829-1878.
Hill, Jennifer, Jane Waldfogel, Jeanne Brooks-Gunn, & Wen-Jui Han. 2005. "Maternal Employment and Child Development: A Fresh Look Using Newer Methods." *Developmental Psychology* 41(6):833-850.
James-Burdumy, S. 2005. "The Effect of Maternal Labor Force Participation on Child Development." *Journal of Labor Economics* 23:177-211.
Johnson, Rucker, Ariel Kalil, & Rachel Dunifon. 2012. "Employment Patterns of Less-Skilled Workers: Links to Children's Behavior and Academic Progress. *Demography* 49:747-772.
Korenman, S. & Robert Kaestner. 2005. "Work-Family Mismatch and Child Health and Development: A Review of the Economics Research." Work, Family, Health and Well-Being, edited by S. Bianchi and L. Casper. Mahwah, NJ: Lawrence Erlbaum. 297-312.

Liu, Qian, & Oskar Nordstrom Skans. 2010. “The Duration of Paid Parental Leave and Children’s Scholastic Performance.” *The B.E. Journal of Economic Analysis and Policy* 10(1):1-33.
Morrissey, Taryn W. 2013. Trajectories of growth in body mass index across childhood: Associations with maternal and paternal employment. *Social Science and Medicine,* 95: 60-68.
Neidell, Matthew J. 2000 “Early Parental Time Investments in Children’s Human Capital Development: Effects of Time in the first Year on Cognitive and Non-cognitive Outcomes.” UCLA Working Paper #806.
Nes, Ragnhild, Lard Johan Hauge, Tom Kornstad, Petter Kristensen, Markus Landolt, Leif Eskedal, Lorentz Irgens & Margaret Vollrath. 2013. “The impact of child behaviour problems on maternal employment: A longitudinal cohort study. *Journal of Family and Economic Issues*, DOI: 10.1007/s10834-013-9378-8
OECD (2008), Sickness, Disability and Work: Breaking the Barriers (Vol. 3): Denmark, Finland, Ireland and the Netherlands, OECD Publishing, Paris.
Powers, Elizabeth. 2001. “New Estimates of the Impact of Child Disability on Maternal Employment”. *The American Economic Review*, 91(2): 135-139.
Rasmussen, Astrid W. 2010. “Increasing the Length of Parents’ Birth-Related Leave: the Effect on Children’s Long-Term Educational Outcomes.” *Labour Economics* 17(1):91-100.
Rege, Mari, Kjetil Telle, & Mark Votruba. 2011. “Parental Job Loss and Children’s School Performance.” *Review of Economic Studies* 78:1462-1489.
Ruhm, Christopher J. 2000. “Parental Leave and Child Health.” *Journal of Health Economics* 19:931-960.
Ruhm, Christopher J. 2004. “Parental Employment and Child Cognitive Development.” *Journal of Human Resources* 39(1):156-192.
Ruhm, Christopher J. 2008. “Maternal Employment and Adolescent Development.” *Labour Economics* 15:958-983.
Ruhm, Christoper J. 2011. “Policies to Assist Parents With Young Children.” *The Future of Children* 21(2):37-68.
Smolensky, E., & J.A. Gootman. 2003. Working Families and Growing Kids: Caring for Children and Adolescents. Washington, DC: National Academy Press.
Statistics, Denmark. 2011. "Women and men." *Kvinder og mænd.*Ziol-Guest, Kathleen M., Rachel E. Dunifon, & Ariel Kalil. (2013). “Parental Employment and Children’s Body Weight: Mothers, Others, and Mechanisms”, *Social Science and Medicine,* 95: 52-59.

Figure 1: Maternal Employment during a Child's 1st 15 Years of Life

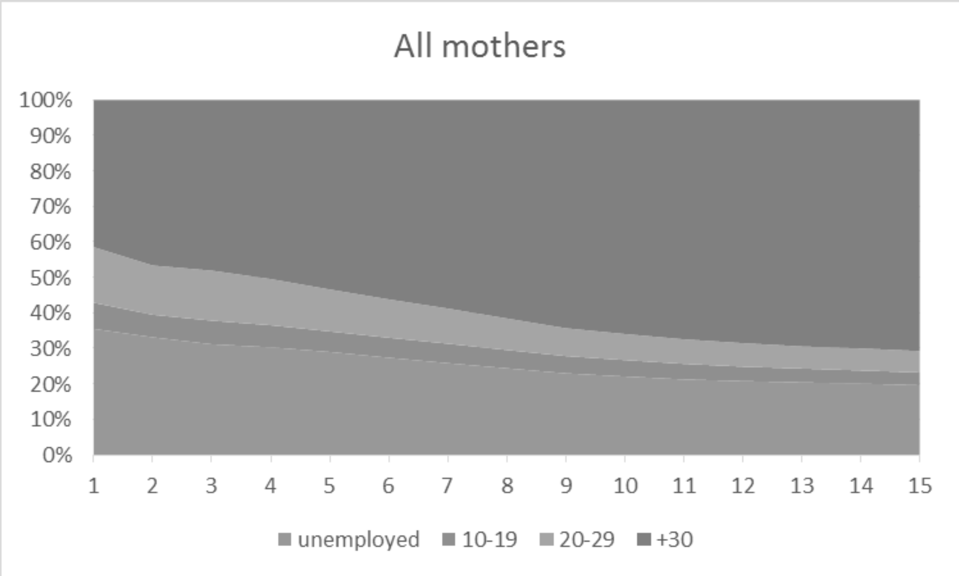

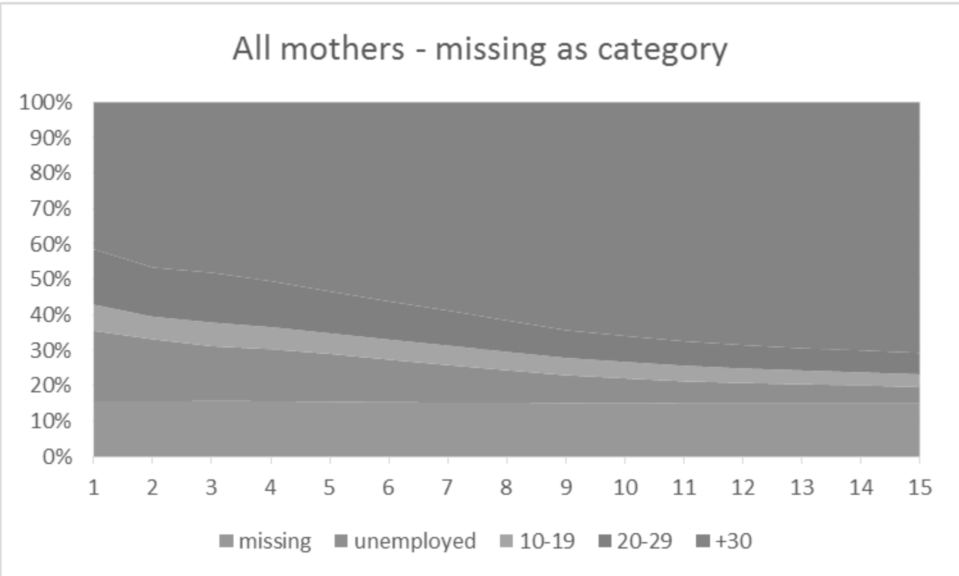

Notes: The figure is the same as the one above, the only difference being that the missing category on employment status is explicitly given.

Table 1: Maternal Employment Patterns during a Child's 1st 15 Years of Life

Percentage of Women in the Sample

Years	0 to 9 hours	10 to 19 Hours	20 to 29 Hours	30+ Hours
0	38.35	57.8	38.31	10.88
1	8.99	23.24	24.37	3.69
2	7.66	10.42	16.26	3.4
3	6.52	4.61	9.75	3.38
4	5.62	2.02	5.29	3.48
5	4.65	0.88	2.78	3.62
6	3.94	0.46	1.42	3.93
7	3.29	0.22	0.74	4.41
8	2.94	0.13	0.43	4.74
9	2.63	0.08	0.26	5.32
10	2.33	0.06	0.15	5.8
11	2.14	0.03	0.11	6.08
12	1.97	0.02	0.06	6.58
13	1.84	0.01	0.04	7.45
14	1.98	0.01	0.03	9.29
15	5.24	0.01	0.01	17.97

Notes: this table reports the percentage of women in the sample by the number of years (out of 15) they spent in a particular work-hour category. For example, 38.35 percent of the women never worked less than 10 hours a week in any of the 15 years.

Table 2: Sample Statistics by Level of Maternal Employment

	Mother's Avg. Work Hours Per Week, Child's 1st 3 Years									
	0-9 hours		10-19 hours		20-29 hours		30+ hours		TOTAL	
Observations	91,628		41,070		62,035		140,466		335,199	
	Mean	SE	Mean	SE	Mean	SE	Mean	SE	Mean	SE
Children										
Standardized GPA at age 15	-0.29	1.01	-0.11	1.00	0.06	0.96	0.19	0.92	0.00	1.00
Female	0.51	0.50	0.51	0.50	0.50	0.50	0.49	0.50	0.50	0.50
Low birth weight (<2500 g)	0.04	0.21	0.05	0.21	0.05	0.21	0.04	0.19	0.04	0.20
Born prematurely	0.05	0.21	0.05	0.21	0.05	0.22	0.04	0.20	0.05	0.21
Chronically ill	0.05	0.22	0.06	0.24	0.06	0.23	0.05	0.22	0.05	0.23
Child born in Denmark by non-native parents	0.12	0.33	0.05	0.21	0.02	0.13	0.01	0.08	0.05	0.21
Birth order	1.09	1.01	1.73	0.88	1.77	0.86	1.68	0.77	1.78	0.92
Mothers										
Elementary school education	0.39	0.49	0.42	0.49	0.30	0.46	0.18	0.39	0.29	0.45
High school education	0.08	0.28	0.12	0.33	0.11	0.31	0.09	0.28	0.09	0.29
Vocational school	0.20	0.40	0.29	0.46	0.33	0.47	0.37	0.48	0.31	0.46
Short-term further education	0.02	0.14	0.02	0.15	0.03	0.18	0.04	0.20	0.03	0.18
Bachelor's degree	0.04	0.20	0.08	0.26	0.17	0.37	0.26	0.44	0.16	0.37
Master's degree or higher	0.02	0.13	0.02	0.15	0.03	0.16	0.05	0.22	0.03	0.18
Missing mothers education	0.24	0.43	0.05	0.21	0.03	0.17	0.01	0.10	0.08	0.27
Age when child was born	26.87	5.20	26.73	4.77	27.89	4.65	28.85	4.20	27.87	4.73
Residing in capital	0.24	0.43	0.26	0.44	0.27	0.45	0.35	0.48	0.29	0.46
Residing in a city	0.44	0.50	0.55	0.50	0.55	0.50	0.51	0.50	0.50	0.50
Residing in a town	0.09	0.29	0.12	0.33	0.12	0.33	0.11	0.31	0.11	0.31
Residing in the outskirts	0.03	0.18	0.05	0.21	0.04	0.20	0.03	0.17	0.04	0.19
Missing geography	0.19	0.39	0.02	0.14	0.02	0.12	0.01	0.08	0.06	0.24
Mothers studying three first years	0.04	0.16	0.08	0.23	0.06	0.18	0.02	0.12	0.04	0.16
Families										
Household average income (0000 DDK)	10.49	9.89	13.51	7.68	15.46	7.61	17.76	6.79	14.82	8.55
Nuclear family all three first years	0.69	0.43	0.85	0.32	0.91	0.25	0.94	0.21	0.86	0.32

Table 3: Coefficient Estimates for a Mother's Employment During Her Child's 1st 3 Years

	1 OLS	2 OLS	3 OLS	4 Maternal FE	5 Maternal FE
Average hours worked per week			0.0029**	-0.0006	
			(0.0002)	(0.0004)	
Categorical work hours					
% of time working (<10 hours omitted):					
- 10 to 19 hours per week	0.324**	0.135**			0.050**
	(0.012)	(0.011)			(0.022)
- 20 to 29 hours per week	0.368**	0.083**			0.001
	(0.008)	(0.008)			(0.018)
- 30+ hours per week	0.506**	0.094**			-0.025
	(0.005)	(0.005)			(0.015)
Observations	335,200	335,200	335,200	335,200	335,200
R^2	0.08	0.22	0.22	0.09	0.09
Full set of control variables	No	Yes	Yes	Yes	Yes

Notes: The dependent variable is a child's grade point average at age 15. Basic control variables include indicators for birth year and birth month. The full set of control variables also includes child's birth order, child's gender, mother's education, whether the mother immigrated to Denmark, an indicator for whether the child was born prematurely, whether the child was low birth weight (i.e., less than 2,500 grams), whether the child has a chronic health condition, the mother's age when the child was born, the household's income before the birth, the percentage of time a child lived with his/her nuclear family, and municipality fixed effects.
*p<.05 ** p<.01

Table 4: Coefficient Estimates for a Mother's Employment During Her Child's 1st 15 Years

	1 OLS	2 OLS	3 OLS	4 Maternal FE	5 Maternal FE
Average hours worked per week			0.0085**	-0.0007	
			(0.0002)	(0.0012)	
Categorical work hours					
% of time working (<10 hours omitted):					
- 10 to 19 hours per week	0.597**	0.041**			0.134**
	(0.025)	(0.023)			(0.102)
- 20 to 29 hours per week	0.619**	0.302**			-0.014
	(0.016)	(0.015)			(0.079)
- 30+ hours per week	0.709**	0.225**			-0.0101
	(0.006)	(0.007)			(0.052)
Observations	335,199	335,199	335,199	335,199	335,199
R^2	0.09	0.22	0.23	0.10	0.09
Full set of control variables	No	Yes	Yes	Yes	Yes

Notes: The dependent variable is a child's grade point average at age 15. Basic control variables include indicators for birth year and birth month. The full set of control variables also includes child's birth order, child's gender, mother's education, whether the mother immigrated to Denmark, an indicator for whether the child was born prematurely, whether the child was low birth weight (i.e., less than 2,500 grams), whether the child has a chronic health condition, the mother's age when the child was born, the household's income before the birth, the percentage of time a child lived with his/her nuclear family, and municipality fixed effects.
*p<.05 ** p<.01

Table 5: Two-Stage Least Squares Coefficient Estimates and reduced form estimates

	3-year measure		**15-year measure**	
	Reduced form	2SLS	Reduced form	2SLS
Average hours per week worked by a mother		0.016**		0.051**
		(0.002)		(0.004)
R-squared		0.20		0.10
		First stage		**First stage**
Unemployment rate in municipality	-0.0006**	-0.032**	-0.002**	-0.0429**
	(0.0001)	(0.0009)	(0.0001)	(0.0016)
R-squared	0.22	0.25	0.22	0.35
F-statistic		1,180		1,415
Observations	335,199	335,199	335,199	335,199
Full set of controls	Yes	Yes	Yes	Yes

Notes: The dependent variable is a child's grade point average at age 15. Basic control variables include indicators for birth year and birth month. The full set of control variables also includes child's birth order, child's gender, mother's education, whether the mother immigrated to Denmark, an indicator for whether the child was born prematurely, whether the child was low birth weight (i.e., less than 2,500 grams), whether the child has a chronic health condition, the mother's age when the child was born, the household's income before the birth, the percentage of time a child lived with his/her nuclear family, and municipality fixed effects.
*p<.05 ** p<.01

Table 6: Two-Stage Least Squares Coefficient Estimates for a Mother's Employment During Her Child's 1st 3 Years for Different Subsamples

		Estimate	Std. dev.	Observations
Overall		0.016**	0.002	335,199
Education	No education	0.043**	0.012	124,548
	Education	0.013**	0.002	210,651
Gender	Girl	0.014**	0.003	167,996
	Boy	0.017**	0.003	167,203
Birth month	January-June	0.013**	0.003	168,547
	July-December	0.020**	0.003	166,652
Subjects	Math	0.011**	0.002	335,199
	Danish	0.009**	0.002	335,199
	English	0.021**	0.003	335,199
	Science	0.010**	0.002	335,199

Notes: The dependent variable is a child's grade point average at age 15. Basic control variables include indicators for birth year and birth month. The full set of control variables also includes child's birth order, child's gender, mother's education, whether the mother immigrated to Denmark, an indicator for whether the child was born prematurely, whether the child was low birth weight (i.e., less than 2,500 grams), whether the child has a chronic health condition, the mother's age when the child was born, the household's income before the birth, the percentage of time a child lived with his/her nuclear family, and municipality fixed effects.
*p<.05 ** p<.01

Appendix 1: OLS coefficient Estimates on Full Set of Control Variables

Variables	Coefficient	SE
Birth-year indicators		
1987	-0.424	0.005
1988	-0.450	0.005
1989	-0.481	0.005
1990	-0.506	0.005
1991	-0.417	0.005
1992 (omitted)		
Birth-month indicators		
January	-0.020	0.008
February	-0.012	0.008
March	-0.007	0.008
April	-0.016	0.008
May	-0.032	0.008
June	-0.041	0.008
July	-0.046	0.008
August	-0.056	0.008
September	-0.057	0.008
October	-0.045	0.008
November	-0.022	0.008
December (omitted)		
Birth order	-0.155	0.002
Female indicator	0.198	0.003
Mothers' highest education prior to birth		
No education	-0.260	0.005
High school	0.393	0.006
Vocational school (omitted)		
Short-term further education	0.291	0.009
bachelor degree	0.388	0.005
Master degree or higher	0.664	0.009
Missing education	-0.051	0.012
Immigrant parents		
Child born outside Denmark	-0.340	0.018
Child born in Denmark	-0.254	0.011
Age of mother when she gave birth	0.031	0.000
Child has a chronic health condition	-0.076	0.007
Premature birth indicator	0.039	0.009
Missing gestation	0.017	0.022
Low-birth weight indicator (<2,500g)	-0.099	0.010
Missing low birth weight	0.170	0.025
Percentage of child's first 15 years spent in a nuclear family	0.218	0.007

Family type missing	0.102	0.019
Household income prior to birth	0.005	0.001
Missing household income	0.050	0.010
Constant	-0.642	0.015

Appendix 2: Mothers' mental health in high and low unemployment municipalities

Tables A2.1 through A2.3 show the median for three different variables that measure maternal mental health at the time when the child is three months old, when the child is three-four years old, and when the child is seven years old. The tables show shares of scores for mothers living in high unemployment municipalities and in low. The numbers stem from the DALSC sample. Because the instrument is local female unemployment stratified by education, we compare mothers with a given level of education in one municipality to mothers with same level education in another municipality in column one and two.

Table A2.1: Mothers' mental health at child age three months (survey in 1996), stratified by median unemployment rate

	Mothers living in a municipality with unemployment rate above country median	Mothers living in a municipality with unemployment rate equal to or below country median
Anxiety		
Yes, and seen a doctor	1,94	2,1
Yes, but have not seen a doctor	6,68	6,33
No	91,38	91,57
(Postpartum) Depression		
Yes, and seen a doctor	1,08	1,07
Yes, but have not seen a doctor	3,32	3,1
No	95,6	95,83
Did your feeling of being able to cope with everyday problems change after birth from before your pregnancy?		
Not changed	69,44	70,88
Changed for the better	10,26	9,87
Changed for the worse	20,3	19,25

Table A2.2: Mothers' mental health at child age 3.5 years old (survey in 1999), stratified by median unemployment rate

	Mothers living in a municipality with unemployment rate above country median	Mothers living in a municipality with unemployment rate equal to or below country median
Anxiety		
Yes, and seen a doctor	5,69	5,91
Yes, but have not seen a doctor	5,69	4,64
No	88,62	89,44
Depression		
Yes, and seen a doctor	6,5	5,91
Yes, but have not seen a doctor	8,94	5,97
No	84,55	88,12
Problems coping with everyday problems		
Yes, and seen a doctor	5,69	5,1
Yes, but have not seen a doctor	19,51	21,57
No	74,8	73,33

Table A2.3: Mothers' mental health at child age 7 years old (survey in 2003), stratified by median unemployment rate

	Mothers living in a municipality with unemployment rate above country median	Mothers living in a municipality with unemployment rate equal to or below country median
Anxiety		
Yes, and seen a doctor	8,06	8,07
Yes, but have not seen a doctor	1,61	3,1
No	90,32	88,83
Depression		
Yes, and seen a doctor	9,14	9,16
Yes, but have not seen a doctor	4,3	4,08
No	86,56	86,76
Problems coping with everyday problems		
Yes, and seen a doctor	9,68	9,27
Yes, but have not seen a doctor	19,89	17,6
No	70,43	73,13

For each question in Tables A2.1-A2.3 a dummy takes the value zero for positive answers, such as *not* experiencing anxiety, and 1 otherwise. In Table A2.4 these dummies are the outcomes in each of the nine logit models. The dependent variable of interest is a dummy for mothers living in a municipality with unemployment rates (stratified by education) above the median. Table A2.4

presents the association between local female unemployment rates (stratified by education) and maternal mental health.

None of the estimates in Table A2.4 are significant at either ten or five percent level. These results indicate that the instrument is not weak. The results also show that local female unemployment by education affects the individual choice of working locally, and one aspect that could negatively affect maternal caregiving, poor mental health, are *not* significantly correlated with their location.

Table A2.4: Logit estimates for mothers living in above median unemployment rate, stratified on education, and maternal mental health.

	Logit estimates for Mothers living in a municipality with unemployment rate above country median		
Maternal mental health outcomes	Child age 3 months	Child age 3,5 years	Child age 7 years
Hard coping with everyday problems	0.029	0.001	0.171
	(0.099)	(0.092)	(0.123)
Depression	-0.224	0.165	0.052
	(0.157)	(0.129)	(0.119)
Anxiety	0.161	0.105	0.074
	(0.207)	(0.148)	(0.086)

Notes: The dependent maternal mental health variables: Coping with everyday problems, depressen and anxiety. All models include control variables: indicators for birth year and birth month, child's birth order, child's gender, mother's education, whether the mother immigrated to Denmark, an indicator for whether the child was born prematurely, whether the child was low birth weight (i.e., less than 2,500 grams), whether the child has a chronic health condition, the mother's age when the child was born.
´p<.10 *p<.05 ** p<.01

Appendix 3: Complier Analysis

Angrist and Pischke (2009) present one way of counting the compliers for the instrument. Different instruments can potentially have different subpopulations, even though they generate similar average causal effect. Using Bayes rule, Angrist and Pischke (2009) measure the complier population as:

$$P[D_{1i} > D_{0i}|D_i = 1] = \frac{P[Z_i=1](E[D_i|Z_i=1]-E[D_i|Z_i=0])}{P[D_i=1]} \quad \text{(A3.1)}$$

$P[D_i = 1]$ is the proportion of mothers in the sample who work more than 10 hours per week on average in the first three years. By construction $P[Z_i = 1]$ equals close to 0.25. It measures the probability that mothers, stratified by education, are in the bottom quartile of unemployment, i.e. experience low unemployment. Angrist and Pischke (2009) refer to this as the probability that the instrument is switched on. The last parameter $(E[D_i|Z_i = 1] - E[D_i|Z_i = 0])$ measures the probability of mothers working more than 10 hours per week *and* living in municipalities in the bottom quartile unemployment rate. $P[D_{1i} > D_{0i}|D_i = 1]$ is the proportion of treated that complies, given that the instrument is switched on, i.e. the complier population as percent of the treated. To obtain the complier population as percent of the *untreated*, we divide equation A3.1 with 1- $P[D_i = 1]$ instead of $P[D_i = 1]$.

Table A3.1 displays the results of the complier analysis. The results of interest are the last two columns to the right: The compliance probabilities. Most mothers are unaffected by the local unemployment rate within their educational sphere. They work regardless. Only mothers on the margin, 1.8%, are affected as shown in the table for mothers overall. The complier population as percent of the *treated* ranges from 1.3 % to 4.2 %. They work 10 or more hours due to the low unemployment in their municipalities. The compliers are driven by mothers with no education, in which 4.2 percent comply. It makes perfect sense that mothers with lower education are more affected by their local labor market than mothers with higher educations, who are more likely to search farther. The results in table A3.1 suggests that compliers for mothers with bachelor or more education, above median maternal income, and above median age for mothers, all have 1.3 % compliers. It is likely that all three measures resemble the same underlying mechanism: that educated mothers are older, when they have children, and have a higher income pre-birth, and are less affected by the local demand for labor.

For the non-compliers (mothers working less than 10 hours per week on average for the first three years) 7.9 percent would have increased their work hours, had they lived in low unemployment municipalities. These are compliers as percent of the *non-treated*. Table A3.1 also show that the compliance increase with the level of education. The longer the education for the non-treated, the more likely low local unemployment is to induce taking a job. The decrease over education level reflects that mothers with longer education can easier find jobs than mothers without education.

Table A.3.1 Variables in the complier analysis

Subsample	Sample mean	$P[D=1]$	$P[Z_i=1]$	1st stage, $P[D_1 > D_0]$	Ratio of subsample 1st stage to overall sample 1st stage	Complier population	
						as % of treated	as % of untreated
Overall	1.00	0,81	0,23	0,06	N/A	1,8%	7,9%
Elementary school+ missing education	0,38	0,54	0,20	0,12	1,79	4,2%	5,0%
High school + Vocational + Short education	0,43	0,81	0,25	0,06	0,94	1,9%	8,0%
Bachelor + Masters degree or higher	0,19	0,92	0,25	0,05	0,73	1,3%	14,2%
Larger than median maternal income in year prior to birth	0,50	0,86	0,27	0,04	0,63	1,3%	8,2%
Larger than median age for mothers	0,48	0,78	0,24	0,04	0,63	1,3%	4,4%

Subsidizing rural life: Commuter subsidies and labor market behavior

Lisbeth Palmhøj Nielsen

Department of Economics and Business Economics, Aarhus University, Fuglesangs Allé 4, 8210 Aarhus V, Denmark, and The Danish National Centre for Social Research, Herluf Trollesgade 11, 1052 Copenhagen, Denmark

A R T I C L E I N F O

JEL classifications:
H20
H21
H24

Keywords:
Commuter tax allowance
Labor supply
Governmental intervention
MDID
Exact matching

A B S T R A C T

Commuter tax allowance is an expensive subsidy that is part of tax portfolios in most countries, but the existing literature on commuting in general, and commuter tax allowances specifically, is small. This study uses a large-scale municipal reform in Denmark, in which residents in relatively poor municipalities were induced with an extra commuter tax allowance to increase long distance commuting, to estimate an effect of extra commuter tax allowances on unemployment, and wages. The central government introduced and financed the reform, and because only selected outlying municipalities were entitled to the extra tax allowance, I use matching estimators combined with difference-in-difference to infer a plausible causal effect. Exploiting Danish high quality register data with information on the entire population, I find a significant but small effect of the extra commuter tax allowance on distance to work among residents already employed. The extra subsidy has been unsuccessful in promoting more job opportunities for unemployed residents from outlying municipalities.

*E-mail: lisbeth@palmhoej.dk

1. Introduction

Following changes in technology, comptetition and economic conditions firms change their demand for labor. A well-functioning labor market encourages mobility to increase the supply of labor to the firms. The national target is to ensure the most efficient use of the scarce ressources of labor supply. A high degree of mobility in the labor force, both migration and commuting, increases the labor supply in general. The more mobile the population, the fastervacant positions are filled, and with better matches, than otherwise possible with a small, local pool of applicants. A highly centralized labour market demands a higher degree of mobility of the population. The more mobile part of the population migrates to live near the jobs, whereas the less mobile part is left behind in the outlying areas. In recent years depopulation presents a regional concern in which populations in the outlying areasare characterized by a high degree of unemployment, a low degree of mobility, low house prices, and low education levels. One concern is that local populations are left disadvantaged and, as a direct consequence of their residence location, are locked into spatial poverty traps (Felix Weinhardt, 2014). The national target of maintaing a highly mobile labor supply may coincide with the regional target of avoiding depopulation of taxpayers.

People-relocation policies, allocating residents in poor neighborhoods to less segregated neighborhoods, have generally been unsuccessful in decreasing unemployment among the targeted group (Rosenbaum and Zuberi, 2010). New research suggests that people-relocation policies have, however, been successful in helping children escape the adult-life poverty, despite failing to help the adult residents (Chetty, Hendren and Katz, 2015). Place-based policies, supporting targeted geographic areas, have had more success, though often with negative spillover effects on neighboring areas (Givord, Rathelot and Sillard, 2013). Instead of either relocating people to areas with jobs, or creating more jobs in the areas where people live, a third option is to decrease the transportation barriers to jobs, thus expanding the labor markets from available local to more global, an option rarely explored in the empirical literature on spatial mismatch.

A place-based commuter allowance is one way to expand a local labor market. It compensates commuters for the time and money they spend on commuting to work. By ensuring that only residents in certain regions far from larger labor markets are eligible for the allowance, workers living far from large labor markets have an incentive to travel longer to reach (higher-quality) jobs. Furthermore, if the allowance is place-based, an increase in commuter tax allowance does not infer negative spillover effects per se. Residents of other areas that decide to move to the targeted areas, to receive the extra allowance, benefit from that allowance only if they are employed. Thus the policy only attracts employed individuals who benefit the area through income taxes. Despite the different end goals nationally and regionally, the commuter tax allowance will satisfy both. Nationally the allocation of jobs will be more efficient as the pool of applicants increase. Regionally the allowance will maintain and attract residents with jobs, and potentially increase employment locally.

Frank (1978) points out that job search is not necessarily an individual choice, but one that a family chooses. Increasing commuter tax allowance in one area may increase the labor market attachment of one spouse but not the other (i.e., a tied mover), explaining why negative spillovers from other localities may still occur. For families already living in the area due to optimized job opportunities for one spouse, but not the other (i.e. a tied stayer), the increased tax allowance may cause these tied stayers to search more broadly for better job opportunities.

Given that commuter tax allowance is already part of tax policies of most countries (Harding, 2014) and that governmental spending on commuter allowance is sizable –Denmark spent 3.1 billion DDK (0.5 billion US dollars) on commuter tax allowance in 2002 (The Danish Ministry of Taxation), which is 0.2 percent of the Danish GDP (Statistics Denmark) the lack of knowledge about the effects of subsidizing commuting is surprising. Estimating causal effects of barriers to the labor market has proven difficult because of obvious problems with correlated unobservables: Individuals who live in areas that are exposed to higher commuting allowances are more likely to be unemployed in the first place. The literature on commuting, spatial mismatch, and commuter tax allowance consists of excellent theoretical models (Wrede, 2001; Borck and Wrede, 2005), theoretical models supported by empirical correlations (Patacchini and Zenou, 2005; White, 1986), empirical causality studies of programs that support different areas (see review of literature by Gobilon and Selod, 2014). Nevertheless, the effect of commuter tax allowance on labor market outcomes has yet to be explored.

This article contributes to the literature by examining the effects of easing spatial barriers to job search. Using a governmental reform to establish the causal identification, this article examines the effect of increasing commuter tax allowance on distance to work, employment, and wages in outlying municipalities. The analyses build on Danish high-quality register data that contains, beyond socioeconomic information, individual-level self-reported commuter tax allowance, distance to work, degree of unemployment, and wages for the entire population from 1995 to 2012.

In 2004 the Danish government introduced an extra commuter tax allowance in designated outlying municipalities for daily commutes exceeding 100 km. Using register data from 1995-2012 and propensity score matching, I match the residents in the non-randomly treated municipalities with residents in adjacent municipalities. The matched residents constitute my control group in a difference-in-difference specification. As the increased commuter tax allowance was non-randomly distributed, the effects I find are average treatment effects on the treated (ATET). More generally, this article explores how the introduction of a program of increased commuter tax allowance affects residents in outlying municipalities.

2. Background

2.1 The reform

In May 2003, as part of its regional growth plan, the Danish government proposed an increased commuter tax allowance in a number of outlying municipalities. In the proposal, 44 (of a total 271) Danish municipalities characterized by low economic growth and low relocation, long-distance commuters would have the opportunity to receive a higher commuter tax allowance. The proposal was approved by parliament in November 2003, with 50 municipalities. These municipalities were selected on the basis of special needs (many different issues) and low earned income. The reform took effect in January 2004. As Figure 1 shows, the treated municipalities (grey) were mainly clustered in poor regions with long commuting distances to large labor markets (circles and rectangles).

Figure 1: The 50 treated municipalities and larger labor markets

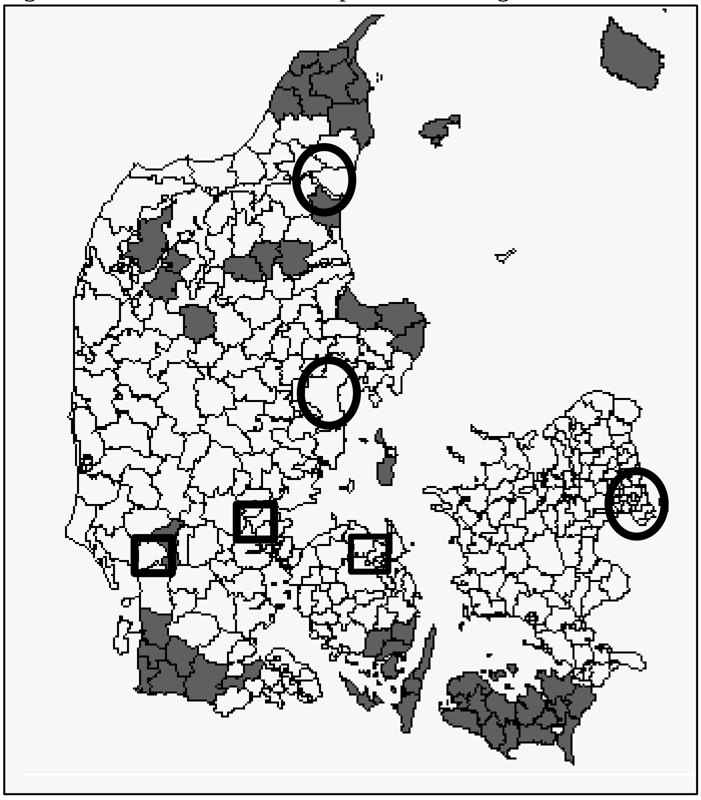

Source: The Danish Ministry of Taxation.

Notes: The circled municipalities employed more than 100,000 individuals in 2013, and the squared employed 50,000-100,000 individuals, whereas the remaining employed less than 50,000 people. The grey are the treated municipalities.[1]

For the 50 municipalities, the reform introduced an increased commuter tax allowance for long-distance commuters. Table 1 shows how the reform worked. Before the reform all residents, regardless of municipality, were entitled to the same commuter tax allowance: for the first 24 km daily, no allowance; the next 25-100 km daily, DKK 1.62 ($0.25) per km; and for farther distances than 100 km daily, DKK 0.81 ($0.125) per km.

After the reform, the commuter tax allowance remained unchanged for the first 0-100 km. The change only applied to commutes farther than 100 km/day. For the treated municipalities, the commuter tax allowance above 100 km/day increased to DKK 1.62 per km, whereas for all other municipalities the rate remained at DKK 0.81 per km, just as before the reform. The rates in Table 1 reflect the tax-deductible amount. For a short introduction to the Danish commuter tax allowance system, see Appendix A.

[1] The treated municipalities are: Arden, Bornholms Regionskommune, Bredebro, Egebjerg, Fjends, Frederikshavn, Grenå, Gudme, Hirtshals, Hjørring, Holeby, Holsted, Højer, Højreby, Lundtoft, Læsø, Løgumkloster, Løkken-Vrå, Maribo, Marstal, Morsø, Nakskov, Nr. Alslev, Nr. Djurs, Nykøbing-Falster, Nysted, Nørager, Ravnsborg, Rougsø, Rudbjerg, Rudkøbing, Rødby, Sakskøbing, Sallingsund, Samsø, Sejlflod, Sindal, Skagen, Skærbæk, Spøttrup, Stubbekøbing, Svendborg, Sydfalster, Sydlangeland, Sæby, Tinglev, Tranekær, Tønder, Ærøskøbing, Aalestrup.

Table 1: The commuter tax allowance before and after the reform in DKK/km (2004 prices) for treated and non-treated municipalities.

	Before	After	
Commuting distance	*All municipalities*	*Non-treated*	*Treated*
0-24 km/day	0	0	0
25-100 km/day	1.62	1.62	1.62
More than 100 km/day	0.81	0.81	**1.62**

Notes: **Bold** indicates the change in commuter tax allowance after the reform

2.1.1 Impact of the reform

Table 2 illustrates the potential impact of the reform on the individual income earner. As 30 percent of the total commuter tax allowance is tax-deductible, Table 2 calculates the tax allowance before and after the reform and the actual net income after the deduction. Table 2 shows that up to the limit of 100 km/day the commuter tax allowance is the same for all municipalities. The extra tax allowance sets in only when the income earner exceeds the limit of 50 km to work (100 km/day).

Table 2: Annual and the total impact of commuter tax allowance (DKK in 2004 prices)

	Before		After		After – before	
Distance/day	*Total commuter tax allowance*	*Total net earnings from tax allowance*	*Total commuter tax allowance*	*Total net earnings from tax allowance*	*Extra allowance due to the reform*	*Annual net earnings due to the reform*
100 km	25,855	7,757	25,855	7,757	0	0
120 km	29,257	8,777	32,659	9,798	3,402	1,021
140 km	32,659	9,798	39,463	11,839	6,804	2,041
160 km	36,061	10,818	46,267	13,880	10,206	3,062
180 km	39,463	11,839	53,071	15,921	13,608	4,082
200 km	42,865	12,860	59,875	17,963	17,010	5,103

Notes: The calculations take into account that 30 percent of the total commuter tax allowance is tax-deductible, and assume a 210 working days per year.

After the reform, driving 120 kilometers secures the individual extra DKK 9,798 net income per year. In Denmark in 2004 this was about 70 percent of an extra month of disposable income for the average income earner (DKK 14,816 in 2004, Statistics Denmark). Driving 120 km before and after the reform, the extra income amounts to 10.4 percent of the entire allowance before the reform, and more for commuting farther (28.3 percent for commuting 160 per day). Driving 120 km to work both before and after the reform automatically increases the income with net income 1.021 DKK annually. For the average person with a disposable income of 14,817 DDK per month, the change in the commuter tax allowance amounts to a 0.6 percent increase in the net income (1.7 percent at 160 km per day). Whether the change is large enough to envoke residents to apply for more distant jobs is a purely empirical question.

One could imagine that a decrease in for instance the price of petrol would encourage residents to take up jobs farther away, the same way as an extra commuter tax would. The differences between a petrol decrease and an extra commuter tax allowance of this size, is not only the extra disposable income, but also the stability of the subsidy. With the extra commuter tax allowance the residents in that municipality was promised the subsidy for at least five years (in 2006 this was prolonged with another five years), as long as they stayed in that particular municipality. This is not the case for the volatile petrol prices.

2.2 Literature review on the potential mechanisms

Local policies for reducing poverty in poor areas have been implemented in various shapes, and in various countries. Gobilon and Selod (2014) present an overview of the literature on the causes of spatial mismatch, which Haas and Osland (2014) describe as geographical barriers to seeking, obtaining, and keeping jobs, and on how governmental interventions attempted to solve the problem. Three possible strategies emerge: moving people closer to jobs, moving jobs closer to people, or easing the transportation barriers to jobs (Ihlanfeldt and Sjoquist, 1998). Examining the first solution—moving people to jobs—Katz, Kling and Liebman (2001) explored the Baltimore, Maryland, Moving-to-Opportunity program, which helped poor households in segregated areas to move to other non-segregated neighborhoods. They find no effect on the program's labor market outcomes. Comparing the findings of the evaluations of MTO to a similar U.S. program named Gautreux, Rosenbaum and Zuberi (2010) find only small effects of both programs.

Examining the second solution, moving jobs closer to people, Busso and Kline (2008) study the introduction of enterprise zones that provide fiscal incentives for businesses in targeted areas in the U.S. They find that these zones reduced unemployment in those areas. Givord, Rathelot and Sillard (2013) study the French business zones (ZFU) that exempt local businesses from taxes, finding significant effects on both local business creations and employment, albeit with negative spillovers on neighboring neighborhoods. The third solution, easing the transportation barriers to jobs, is the strategy that this paper examines through subsidized commuting.

Kain (1968) was the first to address the complex mechanisms between housing and labor market opportunities. Housing segregation in the U.S. post World War II era was substantial, with African-Americans living in certain city neighborhoods and whites living in the suburbs. As businesses decentralized and relocated to the suburbs, the African-Americans still living in the cities, often with no formal education, were left disadvantaged. Due to the longer distances to the job market and the difficulty of reaching certain jobs, African-American workers were subject to higher costs of seeking employment or discouraged from seeking jobs at all. Furthermore, African-Americans would have less information about jobs far away from themselves and their friends. In newer research, Manning and Petrongolo (2011) support the finding that the cost of long distance to a labor market is high labor markets are local, and consequently residents search for jobs close to home.

While after World War II U.S. businesses moved away from the cities, the past three decades have shown the opposite tendency with increased centralization in larger cities at the expense of the outlying areas. This tendency is worldwide. Consequently, the population with the fewest resources is left with depressed housing and labor markets and a higher cost of job seeking. The Danish commuter tax allowance reform aims at expanding the size of local labor markets by including parts of labor markets farther from home. The reform has the potential to enlarge labor markets for residents in outlying areas without forcing them to move closer to the centralized labor markets.

Labor markets and migration are closely intertwined with commuting. The choice of where to live and where to work may be a joint decision in which the individual decides on both residence and job at the same time or sequentially, in which case the decision of either housing or job is decided before the other. Commuting also serves as a substitute for migration and vice versa. Instead of relocating closer to the labor market and having to pay for expensive housing, individuals may choose to commute. But they may also decide to move closer to the job market instead of commuting. Haas and Osland (2014) find that housing discrimination, relocation of jobs to different areas (e.g., through outsourcing), and a lack of accessible transportation create deprived areas. The outlying municipalities in Denmark also experience outrourcing of firms, less public transportation, and housing segregation (through the difficulty to buy property due to capital requirements), explaining the downward spiral. Reducing the price of commuting in the outlying municipalities, the idea is that the reform not only favors longer commutes but also discourages the local population from moving to other closer-to-the-labor-market municipalities and encourages others to move into the treated municipality. Remebering the relatively modest size of the reform (an extra 0.6 percent annual income for commuting 120 km/day) the reform is only one of many aspects included in the decision-making when residents decide where to live.

Still these mechanisms constitute the effects that politicians anticipate. The reform could potentially affect both the extensive margin (between being employed and unemployed) and the intensive margin (working shorter or longer hours). For the extensive margin, a percentage of residents in the outlying municipalities may find the spatial distance to relevant labor markets too far from their homes. If so, they become unemployed solely due to the distance to a potential job. Both nationally and regionally this is inefficient. As the extra commuter tax allowance subsidizes the extra cost of farther commutes and reduces the spatial obstacle to the labor market, this compensation could lead to unemployed residents' finding jobs by searching farther away from home. Following the reasoning of Kleven, Landais, Saez and Schultz (2013) the reservation wage of the job seeker decreases with a decrease in taxes (or increase in subsidies) allowing employers to lower their wage offers. Because of the commuter tax allowance, the residents might be willing to accept a lower wage than they would without the extra allowance. I will test the hypothesis in the empirical section, but anticipate that the size of the reform may be too small for employers to respond to.

For the intensive margin, residents in the outlying municipalities already working could search for more relevant jobs farther away, supporting a better match between employer and employee. This better match would result in lower wages for those using the allowance to compensate for a better, lower-paid job within their current (still far away) job market. For those using the allowance to expand their job market to a larger geographic area, to find a better match than they would have before the extra allowance, the wages would increase. Residents with a preference for commuting may be encouraged to find jobs commuting even farther.

Furthermore, the extra commuter tax allowance in the treated municipalities could attract working residents outside the treated area, persuading them to relocate to the area. For these residents already holding jobs, the wages would be unaffected by the reform. For these reasons, I examine unemployment and wages in the empirical analyses (Appendix B shows the results of the relocation analysis).

3. Data

This study builds on high-quality Danish register data, maintained by Statistics Denmark and contains information about the entire Danish population for 1995-2012 (some outcomes only to 2010). For each year the registers hold unique identifiers that link to individual behavior.

The most important part of the analyses is to discover whether residents responded to the reform. To do so, I use distance to work and the self-reported commuter tax allowance. If the reform had the intended effect, the data would reveal an increase in distance to work and, as a consequence of longer commutes, a larger self-reported commuter tax allowance on annual taxes. However, a discrepancy exists between distance to work and self-reported commuter tax allowance. The distance to work in the registers states the distance from the center of the residence municipality to the center of the work municipality. The self-reported commuter tax allowance is more accurate as it states the exact, actual number of kilometers driven to and from work. If residents decide to work from home on a given day, the commuter tax allowance on that day should be zero. Yet Paetzold and Winner (2014) show that self-reporting on tax returns may be overestimated. As residents learn more about how the system works and behave strategically, researchers should be careful not to necessarily interpret an increase in self-reported commuter tax allowance as an actual change in behavior. In a Danish study Kleven, Knudsen, Kreiner, Pedersen, and Saez (2011) show that the same is the case in Denmark. Therefore, the best indicator of response to the reform is the distance to work measurement combined with the self-reported tax allowance.

Further examining the responses to the reform, I use the self-reported commuter tax allowance to explore the extensive and intensive take-up margins. For the extensive take-up margin, I use a dummy for self-reported commuter tax allowance with the value of 1 if the allowance is greater than zero and 0 otherwise. An increase would suggest that more people were driving farther to work. As the first 12 km to work are not subsidized, the dummy cannot capture an increase in residents commuting short distances to work. For the intensive take-up margin, I construct a dummy that takes the value 1 for residents that take-up more than DKK 25,855 ($6,130) annually, and 0 otherwise. As Table 2 shows, DKK 25,855 is the threshold for driving farther than 50 km to work. If the number of residents who take up allowance above the threshold increases, it indicates that the residents are responding to the reform

To learn if the reform had an effect, I use two outcome variables: the degree of unemployment measured as a fraction between 0-1000, and hourly wages in DKK. In Appendix B, I also examine relocations. The register data allows me to follow individuals to and from work, and in and out of employment over the years. Apart from the outcomes, the registers also hold information about family types (married, cohabiting with and without children, registered partners, singles) and children at the address. Moreover the data also contains information about the parents' location, allowing me to identify whether parents are located in the same municipalities as their grown children. I use all this information in the subsample analyses.

I include only individuals who are 25-55 years old in 2003, the year prior to the reform. I do so to avoid too much overlap with young student residents, not yet part of the labor force, and older residents changing behavior for pension reasons, and the effects of these life decisions. For all included individuals I have full information about their municipality in 2003.

In the final data set, the treated municipalities hold 239,210 residents in 2003. As Figure 2 shows with different shades of grey, the treated municipalities can be divided into six strata: Northern Jutland, Middle Jutland, Southern Jutland, Fuen, Southern Sealand, and islands.

Figure 2: The 50 treated municipalities divided into six strata: Northern Jutland, Middle Jutland, Southern Jutland, Fuen, Southern Sealand, and islands.

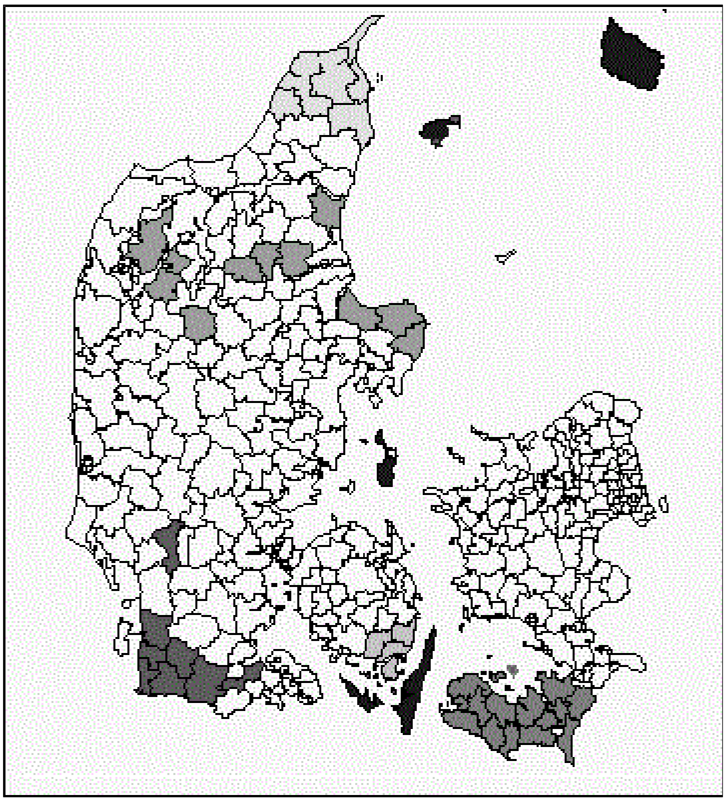

Notes: The shade of grey illustrates each of the six strata.

Table 3 shows that the treated residents are different from the rest of the population. For each stratum, the rows in Table 3 show average unemployment, average disposable and earned income, average number of children per resident, the percentage of residents with children, and the percentage of residents of non-Danish descent. The treated municipalities clearly differ considerably from the others by having poorer residents, who are less likely to relocate, who receive lower wages, and who are more likely to have children. Residents of islands show differences from both the other treated municipalities and the remaining municipalities.

Table 3: Mean characteristics in treated municipalities divided into the six strata: Northern Jutland, Middle Jutland, Southern Jutland, Fuen, Southern Sealand, and islands, and of other municipalities for residents born 1948-1978 in 2003

Characteristics	*Northern Jutland*	*Middle Jutland*	*Southern Jutland*	*Fuen*	*Southern Sealand*	*Islands*	*Other non-treated*
Average unemployment (percent)	2.59	2.09	2.06	2.20	2.08	3.34	2.10
Average disposable income (DKK)	142,498	141,282	141,341	143,238	146,901	136,365	154,063
Average earned income (DKK)	224,081	183,335	181,472	174,167	181,400	207,607	241,249
Average no. of children per resident	0.94	0.98	1.01	0.82	0.80	0.86	0.81
Percentage of residents with children	0.60	0.60	0.60	0.56	0.55	0.56	0.54
Percentage of residents with other than Danish descent	0.04	0.05	0.08	0.07	0.05	0.04	0.09
N	40,371	44,817	27,908	38,299	70,805	17,010	2,176,605

The treated municipalities are located in roughly five different areas of the Danish outlying regions: Northern Jutland, Central Jutland, Southern Jutland, Fuen, and Southern Sealand (see Figure 2). The different areas have varying distances to the nearest large labor market. For example, the nearest large labor market for Southern Sealand (circle in Figure 3) is as far away as Copenhagen (square in Figure 3), approximately 150 km. In 2003, the year before the reform, commuters in non-treated municipalities on average commute 25.8 km, whereas in Southern Denmark the equivalent average is 31.6 km The averages of the remaining treated regions are in the interval 26.2-29.6, all above the non-treated average.

Because of the differences in local labor markets and attitudes towards job search, I expect only the municipalities adjacent to the treated areas to have similar job search possibilites. If, for example, I were to compare the selected municipalities to the capital of Denmark, Copenhagen, the differences in composition would likely include more than we can observe. This makes it impossible to find suitable matches. However, as Figure 1 also reveals, some treated municipalities have no adjacent control municipalities, only other treated municipalities.

As distance to work is critical for obtaining unbiased results, I exclude all treated municipalities with no adjacent control municipality, to ensure that I find the most similar control group possible. These excluded municipalities always have longer commutes than nearby municipalities and differ too much from the treated municipalities. Because islands are special in terms of commuting and other characteristics, no municipalities have similar commuting needs. Moreover, Table 3 shows that the characteristics of treated islands are very different from those of other municipalities, explaining why islands have no obvious control group except other islands. However, as Figure 2 shows, most islands are treated. Therefore, I also exclude all islands.

Figure 3 shows the treated municipalities and the selected adjacent control municipalities. The white spaces in the figure are treated municipalities that are excluded from the analyses, and the dotted rectangles surround the four areas included in the analysis.

Figure 3: Included and excluded treated, and adjacent municipalities.

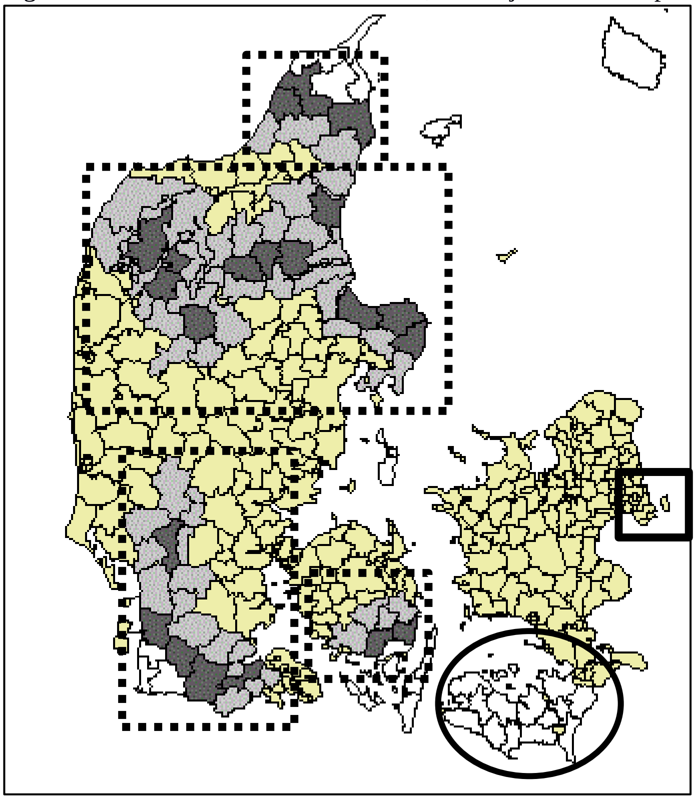

Notes: The white municipalities have been excluded due to missing control municipalities. The dark grey illustrates the included treatment municipalities, the light grey the adjacent control municipalities. The yellow municipalities are the remaining ones. The dotted rectangles display the four regions that are used in the matching. Each treated individual is matched to a non-treated individual in the same region, i.e exact matching. The solid rectangles present Copenhagen and the solid circle, the region of Lolland-Falster that is left out of the estimation.

4. Empirical strategy

Following Lechners' (2002) research design, I use the Danish commuter tax reform as an exogenous shock to the treated municipalities. Using the repeated observations on residents in all municipalities, I compare the difference in outcomes before and after the reform for the residents in the treated municipalities to the difference in outcomes for residents before and after the reform in non-treated municipalities.

Optimally, I would estimate the causal effect of the increased commuter tax allowance on individual i: $\Delta_i = Y_i^1 - Y_i^0$, where Y is a given outcome for individual i, 1 is the treated and 0 is untreated. While I would like to know whether people would have commuted had they *not* been exposed to the reform, and how access to the extra allowance would affect them had they been exposed, the residents cannot be treated and non-treated at the same time, and either Y_i^1 or Y_i^0 is missing. Therefore, because of the missing counterfactual, I use an appropriate control group with similar characteristics to estimate the counterfactual to the treated.

4.1 Individual level fixed effects panel data strategy

To estimate the ATET of increased commuter tax allowance on distance to work, self-reported commuter tax allowance, unemployment and wages, I use a Difference-in-Difference (DiD) strategy combined with matching, similar to Lassen and Serritzlew (2011). DiD is a well-established method for evaluating policy changes (Angrist and Pischke, 2008; Blundell and Costa Dias, 2009; Imbens and Wooldridge, 2009—see Lechner 2011 for a historical overview of the construction of the DiD estimator).

Individual level data produces an excellent opportunity for estimating a fixed effects model on the treated individuals. Equation (1) shows a model that explains how the outcome of individual i before the reform is affected by the treatment D_i, other time-invariant observable characteristics Z_i, and unobserved time-invariant characteristics α_i. D_i assumes the value 1 in the treatment period A, and zero in period B. Equation (2) shows the same model, only changing the subscript B to A—after the reform. By subtracting (2)-(1), I obtain the fixed effects model in which each individual serves as her or his own control group. The time-invariant observed and unobserved characteristics cancel out, and I am left with a reduced model that shows how differences in commuter tax allowance affect the difference in outcomes.

$$Y_{iB} = \theta + D_{iB}\,\alpha + X_{iB}\,\beta + Z_i\,\gamma + \alpha_i + \varepsilon_{iB} \tag{1}$$

$$Y_{iA} = \theta + D_{iA}\,\alpha + X_{iA}\,\beta + Z_i\,\gamma + \alpha_i + \varepsilon_{iA} \tag{2}$$

$$\Delta Y_i = Y_{iA}-Y_{iB} = \alpha\,(D_{iA}-D_{iB}) + \beta\,(X_{iA}-X_{iB}) + (\varepsilon_{iA}-\varepsilon_{iB}) \tag{3}$$

A threat to the identification is the omitted time-varying variables that could potentially change within the individual. Using the reform as exogenous identification, the fixed effects model produces only unbiased results if nothing else changes at the same time as the reform. The time trends are a major factor in evaluating reforms. Changes in the economic environment might, for example, increase the chances of employment in the treatment group but not in the control group. In the fixed effects model, this increase in employment chances would wrongly be assigned to the effect of the reform. As I analyze a longer period in time, similar time trends in treatment and control groups are important.

To take the time trend into account, I compare the differences within the treated individuals to the differences within the control group. After careful matching, the treatment and the control groups differ only in treatment, and equation (4) shows ATET that I measure in the DiD model.

$$\Delta Y = (Y^{T}_{iA}-Y^{T}_{iB})- (Y^{C}_{iA}-Y^{C}_{iB}) \tag{4}$$

As the data contains information on the individual level, I can estimate the DiD model applying individual fixed effects (Blundell and Costa Dias, 2009; Lechner, 2011). The subscript i refer to the individual, as each treated individual is compared to the matched counterpart (also referred to as i) in the estimations.

The DiD model with matching assigns each treated individual to a matched comparison individual, and I estimate the effect of the reform by subtracting the “before” from the “after” estimates, as described in equation (4).

By moving to a treated municipality, residents from the matched comparison group are potentially treated. In a DiD setting, a control group contaminated by treatment can potentially hide a true effect. Even though residents from non-treated municipalities can also be treated by moving to the treated municipalities, they have higher transaction costs, such as costs of moving and of finding a job distant enough. Due to the extra costs from relocating, the model assumes that the residents in the treated municipalities are more likely to be affected by the reform than the non-treated residents, because they already live in the treated municipality. Furthermore, the reform means only a small income increase (less than one percent driving 50 km to work), so only residents at the absolute margin will choose a treated over a non-treated municipality due to the reform alone. If the non-treated residents move to the treated areas to access the extra commuter tax allowance, I expect the effect on distance to work to decrease, but the effect on relocating to increase. Appendix B shows that the relocation pattern does not increase for the matched comparison group after the reform.

4.2 Nonrandom assignment

The most critical threat to the identification strategy is to specify a control group correctly. As the treatment was assigned to outlying municipalities located the farthest from the larger labor markets, and with these municipalities being the poorest, finding control municipalities is not an easy task. Even though the matched comparisons are perfectly matched to the treated on individual characteristics, and thus resembles a random assignment, the treated municipalities should not differ too much on municipal specific characteristics. If they do, the matching will not catch all relevant unobservables. Table 4 shows several municipal economic indicators for treated municipalities, adjacent municipalities bordering the treated municipalities, and the remaining non-treated municipalities in 2003 and 2006, respectively. The t-value state if the change over time is significant.

For identification the trend is the variable of interest. Table 4 shows that the percentage of home ownership and the number of reported burglaries have significantly decreased over time in all three groups. Performing another t-test comparing the trends for the treated and the adjacent municipalities shows that the trends are *not* significantly different. None of the three indicators: percentage of public housing, socioeconomic index, and average class-size, have significantly changed over the period. All in all Table 4 supports the empirical strategy of comparing the treated to the adjacent. Furthermore, Table 4 also shows that the absolute values of most economic indicators are closer in the two than that of the remaining non-treated municipalities.

Table 4: Means of municipal characteristics in 2003 and 2006, in treated and their adjacent municipalities

	2003		2006		2006-2003			
	Mean	SE	Mean	SE	Difference	SE	t-value	
The percentage of home ownership								
Treated	69.8	6.1	66.8	5.8	-3.1	2.4	-3.98	**
Adjacent	68.2	8.1	65.5	7.6	-2.7	2.8	-4.34	**
Other non-treated	66.2	11.5	63.6	10.4	-2.6	3.3	-7.45	**
The percentage of public housing								
Treated	8.3	6.9	8.9	6.5	0.6	2.6	0.72	
Adjacent	9.5	5.9	9.8	5.6	0.2	2.4	0.41	
Other non-treated	11.8	7.9	12.2	7.7	0.4	2.8	1.38	
Socio economic index (higher is worse)								
Treated	0.9	0.3	0.9	0.3	0.0	0.6	0.07	
Adjacent	0.7	0.2	0.7	0.2	0.0	0.4	0.19	
Other non-treated	0.8	0.2	0.8	0.2	0.0	0.5	0.22	
Average class size								
Treated	18.3	1.9	18.6	1.4	0.3	1.3	0.68	
Adjacent	18.4	1.4	18.7	1.7	0.3	1.2	1.17	
Other non-treated	19.4	1.4	19.9	1.6	0.5	1.2	3.59	**
Number of reported burglaries per 1,000 residents								
Treated	49.8	21.3	37.9	12.9	-12.0	4.1	-9.14	**
Adjacent	47.0	19.0	40.3	13.9	-6.7	4.1	-7.46	**
Other non-treated	52.0	20.1	40.9	16.0	-11.1	4.2	-24.85	**

Source: Danish Ministry of Economic Affairs and the Interior.
Notes: ** indicates that the coefficient is significant at 5 % level, and * indicates significance at the 10 % level. The means build on information on the 20 treated and 41 adjacent municipalities included in the analysis. Other non-treated refer to the remaning 180 non-treated municipalities.

4.2.1 Anticipation of the reform

Before the reform took affect in January 2004, it was announced to the public through newspaper articles, and political announcements. The government announced the proposal in May 2003, and in August 15, 2003, a nationwide newspaper brought an article, explaining the new tax allowance to the public, including the names of the treated municipalities.[2] Consequently the public was aware of the new commuting allowance about half a year before it was implemented. For residents searching for new jobs, this new information might have affected their search pattern even before the reform set in.

The actual tax allowance did not change prior to the reform, but the distance to work and unemployment rates might have. If the residents did react prior to the reform, the difference-in-

[2] (http://politiken.dk/oekonomi/ECE65706/stoerre-fradrag-til-pendlere-i-udkantskommunerne/)

difference results will be biased. Figures 4-6 shows the differences in means for unemployment in the treated and the matched comparisons that I include in the estimation (see section on matching). Figure 4 shows that the trends in unemployment in year 2003 are not significantly different in the treated and matched comparisons.

Figure 4: Differences in means between the treated and the matched comparisons in unemployment

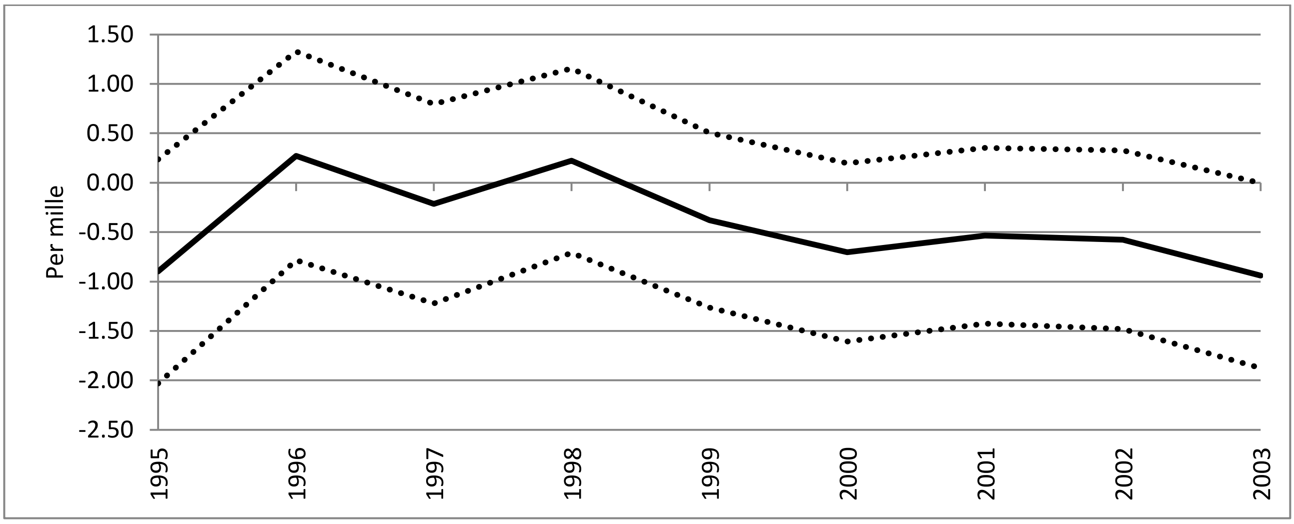

Figure 5: Differences in means between the treated and the matched comparisons in average distance to work

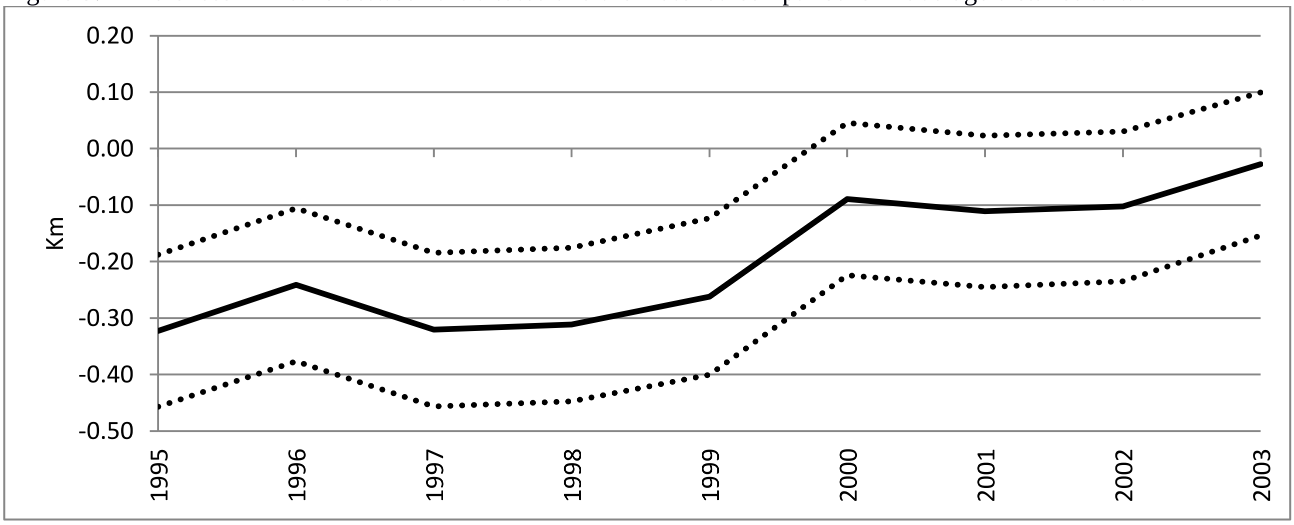

The differences in average distance to work (Figure 5) shows that the trend has been similar in treatment and matched comparison group since 2000. The figure shows an insignificant increase from 2002 to 2003.

If the insignificant increase is caused by the reform, we should also see an increase in the long-distance commuting in 2003. To explore this further, Figure 6 shows the differences in treatment and matched comparisons for residents who have a distance to work farther than zero kilometers or within their own municipality. Figure 6 shows a slight increase from 2002 to 2003, but examining the entire period, this increase is small and insignificant.

Figures 4-6 shows weak signs that could indicate anticipation effect. These are more likely due to yearly fluctuations, as they are small and insignificant. Therefore I continue with the analysis matching the residents in year 2003, and including year 2003 in the estimations of the effect sizes.

Figure 6: Differences in means between the treated and the matched comparisons in distance to work, given above 0.

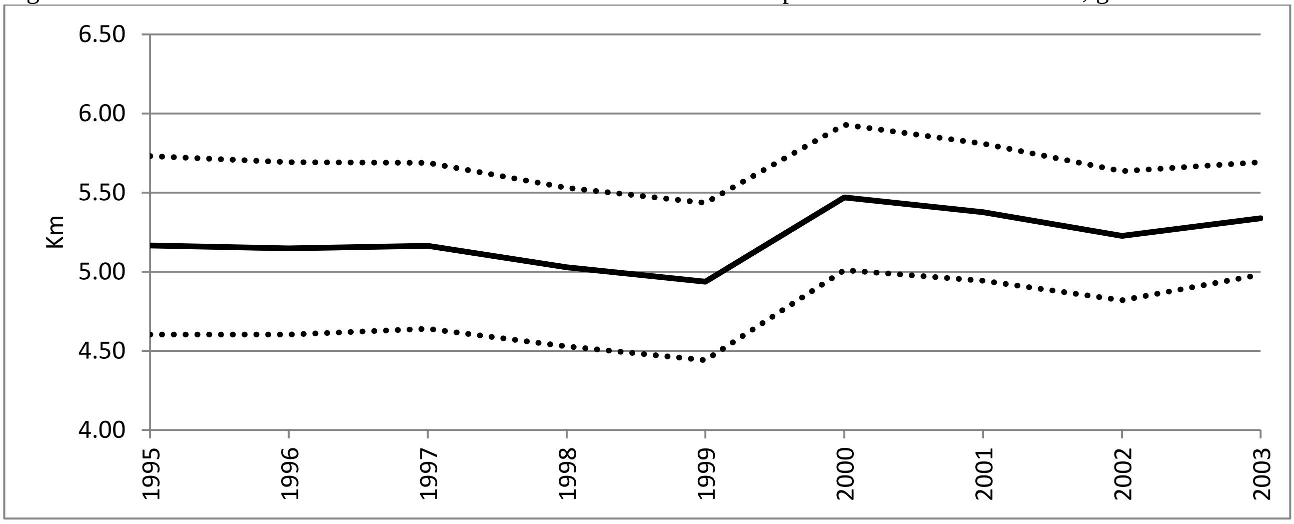

4.3 Conditional Independence assumption and exact matching

Matching is widely used in non-experimental evaluations of policy programs (Heckman, Ichimura and Todd, 1998). It is aimed at reconstructing randomized trial conditions, assigning treatment to a treatment group and comparing treatment to a suitable matched comparison group. If the matching is successful, I can attribute the differences in outcomes between treatment and matched comparisons to the policy program.

The most important assumption in matching is the Conditional Independence Assumption (CIA). It states that observables can fully explain differences in treatment and control groups, i.e., the two groups can differ in composition only for the observed variables. The propensity score matching procedure then correctly re-weights the groups.

4.4 Propensity score matching

Using a dummy for assignment to treatment, and taking the value 1 for residents in treated municipalities in 2003 and zero for adjacent municipalities as the outcome variable, a probit model estimates the propensity scores. The covariates in the model are all observable variables that influence the assignment to treatment and that are necessary to correctly balance the weights. The propensity scores, which assume values between 0 and 1, state how likely the individual is to be assigned to treatment.

Using the estimated propensity scores among the treated municipalities (dark grey in Figure 3), I match individuals resembling an individual in the local control area (light grey in Figure 3). As Lechner (2002) I use the estimated propensities that balance the sample, and complement them with exact matching on region, i.e., I ensure that the matched comparisons are indeed selected among the adjacent municipalities in the same region as the treated municipalities. The dotted rectangles in Figure 3 displays the four regions I use in the matching procedure.

The observables included in the probit estimation are crucial for the matching, as is including all relevant variables that affect how people search for jobs and housing. If some unobservable variable is greatly important to the results, propensity score matching will not make up for the bias in the results unless the unobservable correlates with the observables. Appendix C shows the variables included in the Probit estimation for calculating the propensity scores. As only the mean independence of who is in the treatment group and who is in thematched comparison group matters, the estimates themselves are uninteresting.

Like Leth-Petersen (2010), I begin by matching in the year before the reform (2003). However, to make the match more accurate, I also use covariates in the Probit estimation for the years prior to 2003. I have pruned the matching to include variables only that have predictive power with p-values greater than 0.20. Including variables with poor predictive power increases bias in propensity score matching estimates (Caliendo and Kopeinig 2008). The few variables that have lower predictive values than 0.20 have proven to be important to the balancing of the matching.

I exclude variables from the matiching with p-values abo For unemployment, I match every year from 1995 through 2003. I match the years 2000 through 2003 for most covariates. For some covariates that change less often, such as family type and number of children, I match for 2002 or 2003 only. The weights for each observation (1 for each treated individual and 1 or higher for each selected matched comparison) are saved and assigned to that individual for all observations before and after the reform. This weight is assigned to all estimations and descriptive tables.

I perform nearest neighbor propensity score matching with both caliper (0.000006) and with replacement. The nearest neighbor matching selects the matched comparison with the propensity score closest to each treated individual and connects the two outcomes. Matching with replacement allows reselection of each matched comparison multiple times. By allowing for replacement, the matched comparisons are closely matched, and thus the bias becomes smaller, however, at the cost of a higher variance. One important aspect of matching with replacement is to order the data randomly (Caliendo and Kopeinig, 2008).

By setting the caliper to 0.00006, I limit how distant the match can be. If the data set does not hold any individual that resembles the treated individual closely enough, the treated individual retains no weight and drops out, as do the control individuals without a suitable match. With the caliper at 0.000006 the number of treated individuals in the analyses drops from 94,191 to 82,590 (12.3 percent decrease). There is a trade off between balancing the observations, in this case on not only one facet, such as unemployment, but also on multiple other facets, such as distance to work, annually wages and relocation. The size of the caliper is as high as it can be while still balance the averages as closely as possible.

Appendix D shows the Kernel distrubitions of the treated and matched controls before and after the matching in each of the four regions. The distributions before the matcing follow the same pattern in the treatment and control group, but the balancing tests, also in appendix D, shows that the averages are quite different in the two groups. The balancing tests show variables that potentially could impact the effect estimates. For each variable appendix D shows the means before and after matching in treatment and matched comparison groups, the percentage bias, the percentage reduction in bias, and the corresponding t-values that proclaims whether the after matching means are significantly different. A bold t-value indicates that the means of the treatment and the matched comparisons are *not* the same. The distributions after the matching are close to the before-matching ones, but with the exeption that more matches have been made around the center and less from the tales, due to the low caliper.

The balancing tests show that most of the variable means in the two groups are *not* significant differencent after the matching, The matching appears to have produced a comparison group that largely resembles the treatment group. As the means in the natural logarithm of the annual wages in years 2000-2003 differ, wages are less accurately specified, and one must be careful in the interpretation of effects on wages. As I will show in the section of common trends, the trends in wages are similar though in the treatment and matched comparison groups.

4.4.1 Common trends assumption

In this section, I show how the assumption of common trends in important variables are generally met, and the DiD estimations therefore can be causally interpreted.

For any of the previous results to be valid, the crucial assumption for the DiD estimator of common trends in treatment and matched comparisons must be met. Table 5 confirms that the treatment and comparisons are closely matched, whereas the common trend assumption ensures that the treatment and matched comparisons have common *trends* in outcomes before the reform. That is, in both outcomes, and other variables that may be important for the outcomes. The assumption implies that the treatment group, had it not been treated, would have followed the same trends as the group of matched comparisons. If the assumption is met, deviations in outcomes after the reform can be interpreted causally as a result of the treatment. The logic being that because the treatment and matched comparisons have identical pre-reform trends, the resulting effects are most likely *not* due to differences in other characteristics.

Figures 7, 8, 9, 10, and 11 illustrate the differences in how the residents respond to the extra commuter tax allowance—specifically the average distance to work, the average self-reported commuter tax allowance the average distance to work if it is more than zero, and the percentage of residents that report commuter tax allowance for the treatment and matched comparison groups, respectively. If the difference is within the 95 percent confidence interval and the conficence interval overlaps with zero, the differences intrends are *not* significantly different. If they are significantly different, the differences in trends should be constant for a longer period before the reform. For all five figures, the common trends before the reform appear close to identical for the treatment and control groups since 2000. I only use the years 2001-2003 in the estimation. For commuters that take up more than DKK 25,855 (the threshold for commuting farther than 50 km to work) in Figure 12, the percentage has been higher in the treated municipalities, but decreasing since 2000. As this is the affected margin, the effect estimates on distance to work are likely to produce a lower bound estimate, producing a conservative estimate. This is preferred to overestimating an effect.

From all six figures the pre-trends seem to be closely related among the treated and their matched comparisons.

Figure 7: Differences in means between the treated and the matched comparisons in average distance to work.

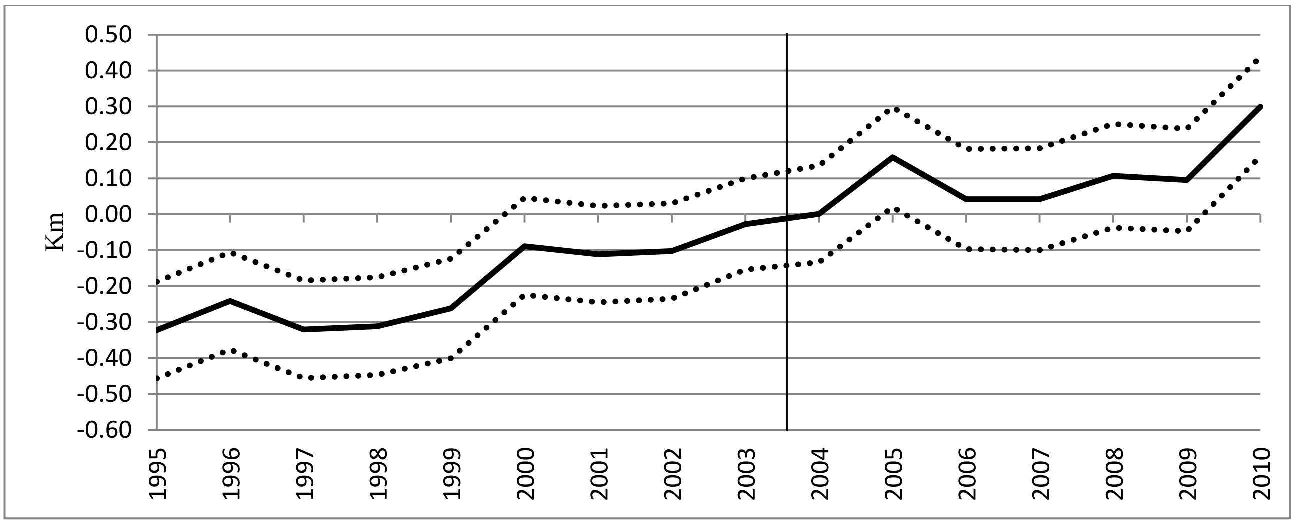

Figure 8: Differences in means between the treated and the matched comparisons in distance to work, given a distance greater than 0.

Figure 9: Differences in means between the treated and the matched comparisons in average commuter tax allowance.

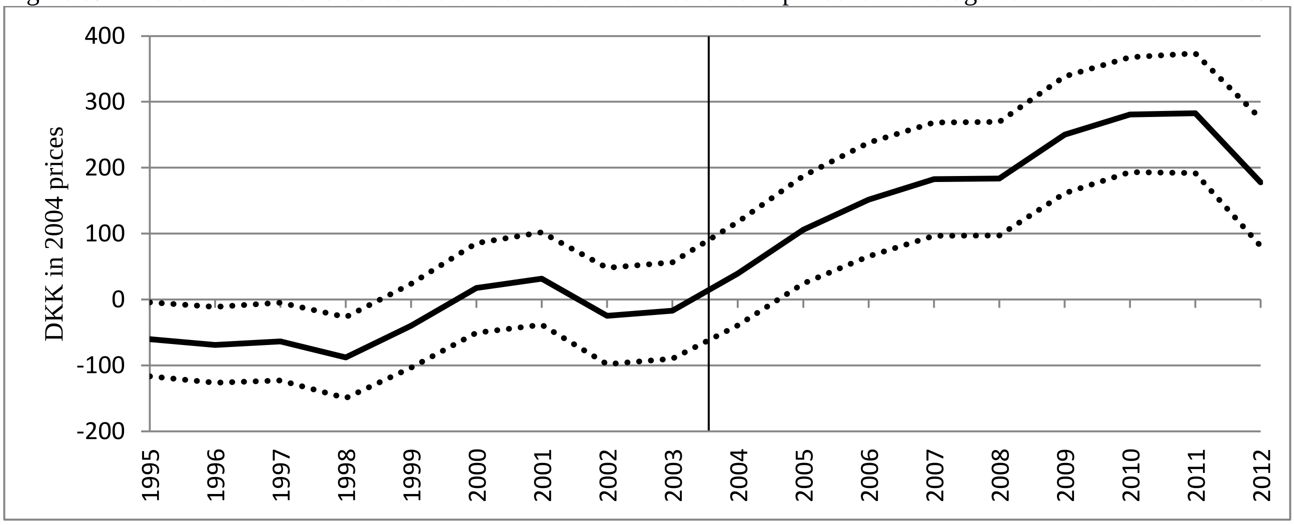

Figure 10: Differences in means between the treated and the matched comparisons in average commuter tax allowance, given a take-up greater than DKK 0.

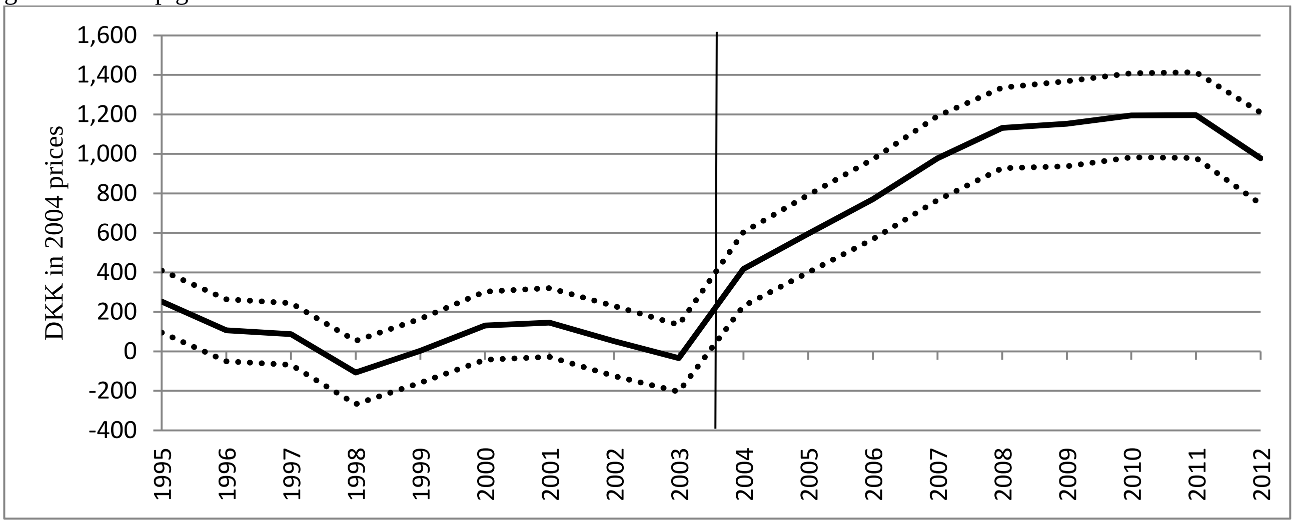

Figure 11: Differences in means between the treated and the matched comparisons in percentage of residents that take up commuter tax allowance.

Figure 12: Differences in means between the treated and the matched comparisons in percentage of residents that take up more than DKK 25,855 (the cut-off) of commuter tax allowance.

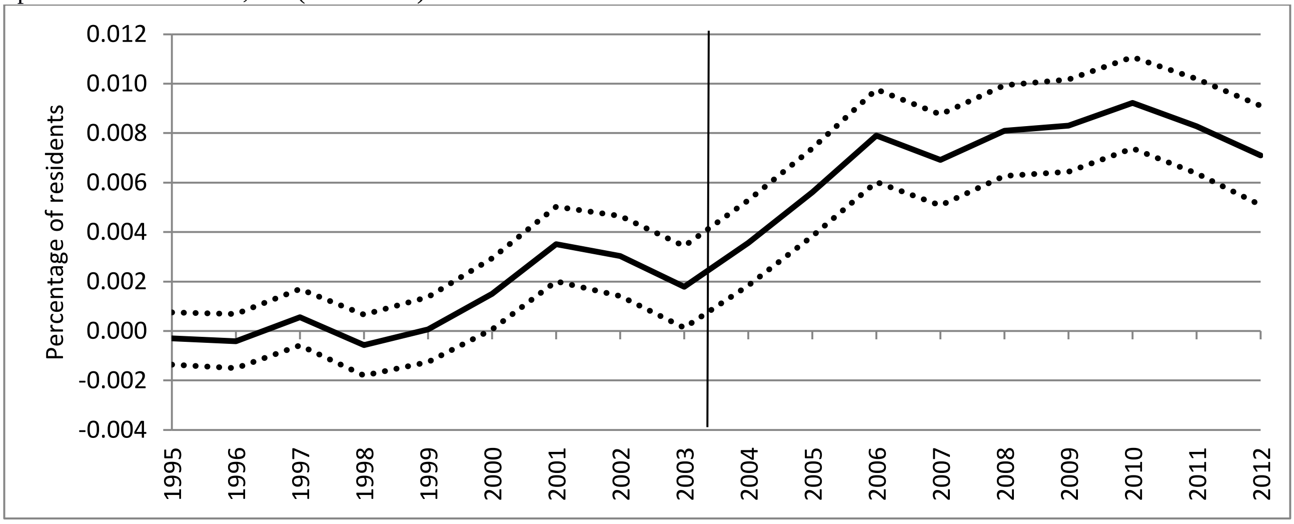

4.4.2 Common trends in outcomes

I explore if the common before reform trends are similar in the treatment and matched comparison groups. Figures 13 and 14 show the differences in trends in the two outcomes: unemployment and wages. For unemployment the trends look similar for the entire period.. Since 2000 the treated have higher log(annual wages) than their matched comparisons, but the trend, however, is similar for both groups from 2000 to 2003. As common trends depend on similar trends, not similar absolute levels, and high as well as low income earners are likely to be affected alike by the reform, the common trend assumption seems to be met.

Figures 15 and 16 display the differences in common trends for two measures that could potentially influence the outcomes. Figure 15 illustrates the trends in percentage of residents that have children living at home. Since 2000 the trends are closely related in the two groups, but the residents in the treated municipalities significantly fewer. Figure 16 shows the differences in percentage of residents with a partner. Even though the matched comparisons are more likely to have a partner at any point before the reform, the trends appear to have the same pattern, except for the year 2001.

Figure 13: Differences in means between the treated and the matched comparisons in unemployment (1000^{th}).

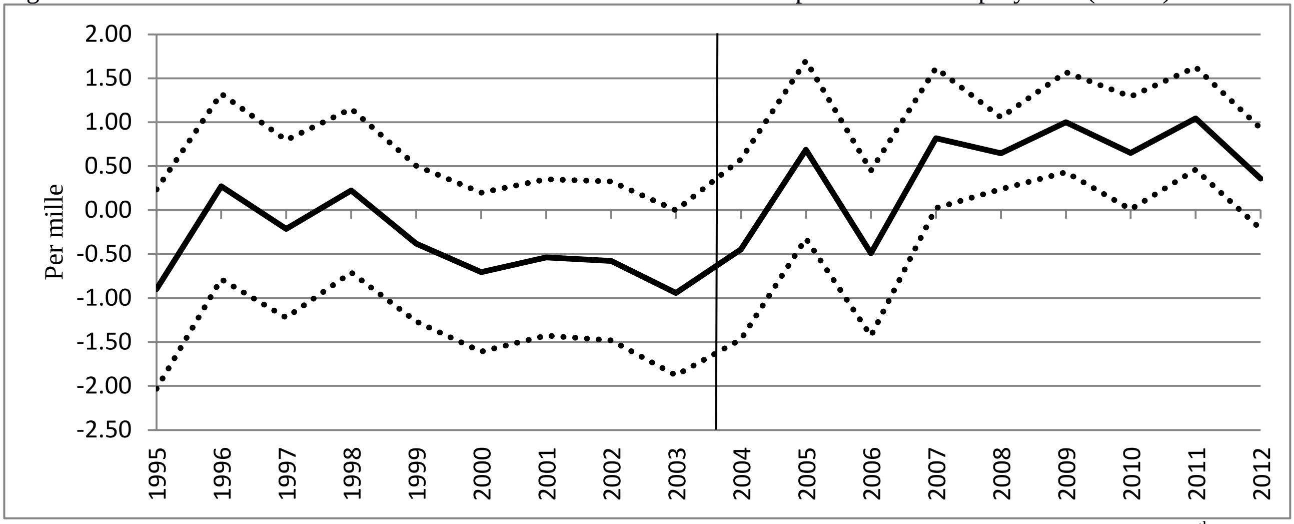

Notes: degree of unemployment measured as a fraction of the year unemployed between 0-1000, denoted (1000^{th})

Figure 14: Differences in means between the treated and the matched comparisons in log(annual wages) in 2004 DKK prices.

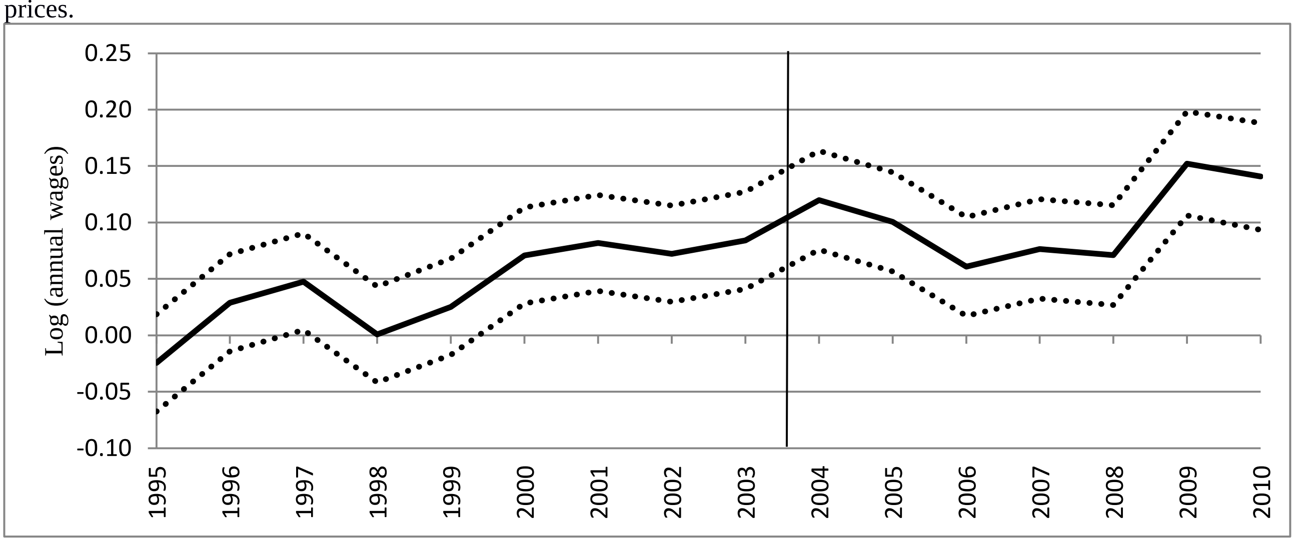

Figure 15: Differences in means between the treated and the matched comparisons in percentage of residents with children living at home.

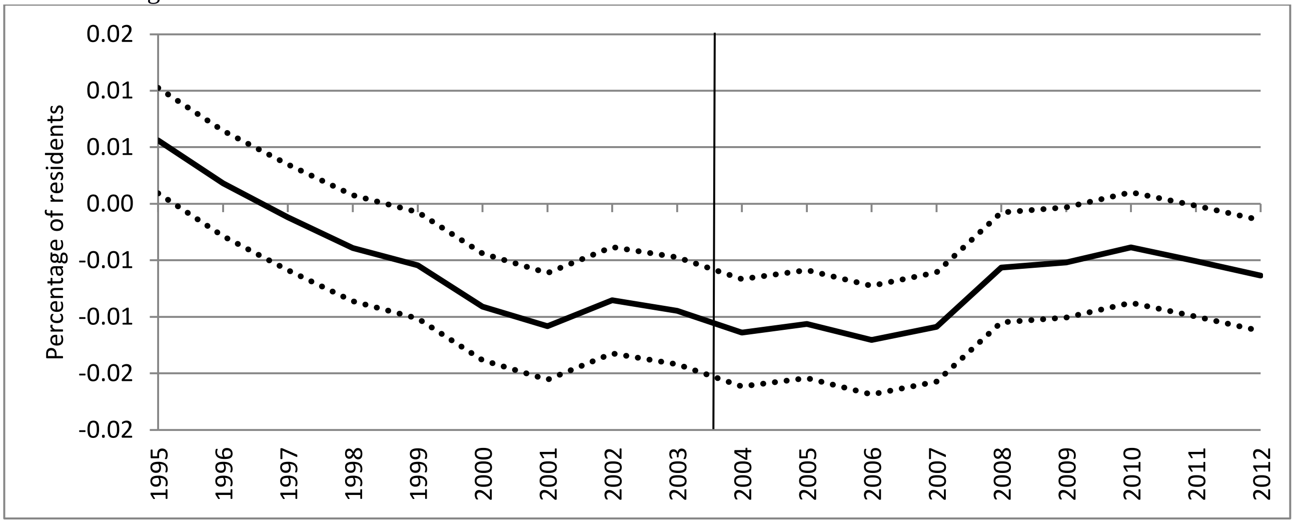

Figure 16: Differences in means between the treated and the matched comparisons in percentage of residents living with a partner.

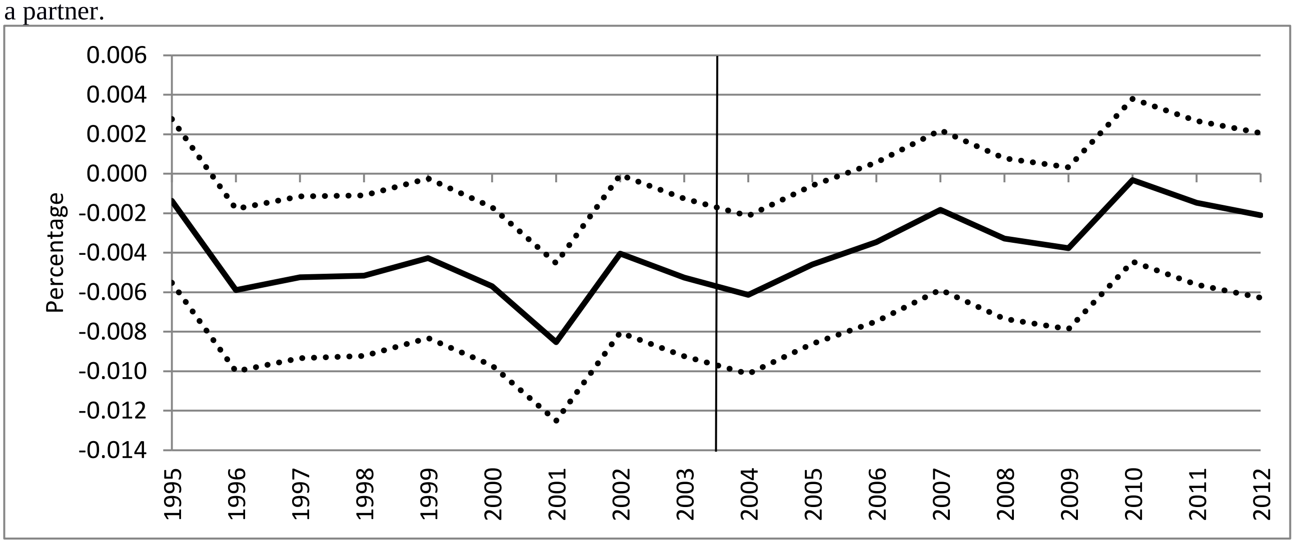

The validity of the results relies on the common trends assumption. The shown trends are very similar both in the two groups, only with a few bumps, and all these analyses are consistent with the notion of common trends. As I cannot exhaust all possible confounders, the common trends shown are reducing the probability of any meaningful bias remaining, in support of unbiased estimates.

4.6 Standard errors on estimated propensity scores

The literature has long been struggling with how to adjust the standard errors when the propensity score estimator itself is estimated (Abadie and Imbens, 2006, 2008). Without a better tool for adjusting the misspecified standard errors, several papers adjust the standard errors by bootstrapping (among others Lassen and Serritzlew, 2011) and I do too. Doing so, however, one paper argues that bootstrapping does not adjust for the misspecified standard errors in the first place (Abadie and Imbens, 2008).

5. Results

5.1 Did the residents respond to the reform?

All residents working farther than 50 km from home automatically receive a higher commuter tax allowance after the reform without changing behavior, *if* they know about the change and report it on their tax returns. As Figures 7-12 have shown, the treated residents increased their take-up in 2004, supporting that they knew about the reform. Table 6 shows the size of the effects in Figures 7 through 12. The first row shows the effect of the reform comparing the average of the three pre-reform years (2001-2003) to the average of the three post-reform years (2004-2006).[3] The following rows present the effects of the reform year to year, comparing each year to the pre-reform average.

In 2007 Denmark experienced a municipal reform, and in 2008 a financial crisis, which could impact the treatment group's unemployment differently from the matched comparison group, counterfactual to what the reform intended. Consequently, the strictly unbiased results of the reform are in the years 2004-2006, whereas the later results may be attributed to different circumstances in the two groups that are unrelated to the reform. Consequently I refrain from presenting the results for years after 2006 in Table 6-10.

The first column in Table 6 displays the effects on distance to work, which increased significantly after the reform. The estimate shows an increase of 0.160 km, compared to the mean of 5.5 km in 2003, is 2.9 percent.[4] All three years after the reform, the estimates are positive with the main influence of 0.22 km in 2005. The reform would affect the residents already driving farther than 50 km to work, and those just below the 50 km threshold. Both groups already have a taste for commuting, and after the reform they have an incentive to commute more often or find a job farther away from home.

The second column in Table 6 shows that the treated residents report DKK 114.12 DKK more than residents in the control group, and the estimate is significant at a 5 percent level. In the first year after the introduction of the reform (2004), the residents from the treated municipalities report DKK 33.67 more than the residents in the control municipalities. In the

[3] Bertrand, Duflo and Mullainathan (2004) show that averages produce more conservative estimates than simply comparing one year pre-reform means to the year post-reform, because the averages take the year-to-year volatility into account.

[4] The 2006 estimate is significant at 10.3 level.

following years, the take-up increases yearly, suggesting that the treated residents commute farther to work than before the reform and that they have full information about the commuter tax allowance.

In year 2006 the distance to work increase less than the year before, whereas the commuter tax allowance increase more than the previous year. The differences could be explained by a higher degree of commuting, even if the work place is the same. The self-reported commuter tax allowance is reported per day. The theoretical foundation as to why residents would increase their commuting to work is less clear. The differences are more likely due to the learning among commuters, as explained previously. When residents have to self-report they rarely lower their reports from one year to the next, and the more they learn, the more the are likely to cheat (Paetzold and Winner,2014; Kleven, Knudsen, Kreiner, Pedersen, and Saez, 2011). I only use these estimates for self-reported allowance to show that the reform was implemented, whereas the more behavioral labor market effects lies within the distance to work estimates.

The third column in Table 6 presents the effect on the probability of using commuter tax allowance. Although the percentage of residents that use the commuter tax allowance decreased after the reform (Figure 7), the extensive margin of no-allowance/allowance is not the margin that the reform was intended to affect. As far as the reform is inducing residents to expand their work distance from a few km to work (residents reporting zero allowance) to long-distance commuting, the allowance has not been successful.

The fourth column in Table 6, however, shows the margin that should be affected by the reform, the commuter tax allowance above DKK 25,855, which is the threshold for commuting 50 km to work daily. From the pool of residents already commuting longer than 50 km to work before the reform, more commute farther than 50 km to work after the reform. The results in Table 6 show a 0.3 percentage point increase on this margin—more treated residents than residents in the control group commute more than 50 km to work—and the results are significant. Of the 2.9 percentage that was above this margin in 2003, this is an increase of 10.3 percent (in 2006 17.1 percent). Although within a small sample of residents, this is quite a large increase.

For residents on the extensive margin, the compensation has not been large enough to trade time for commuting. For residents on the intensive margin, the allowance has changed their behavior significantly. As a direct result of the reform, residents do commute farther and take up more commuter tax allowance.

Table 6: Average treatment effect on the treated: Distance to work in km, average commuter tax allowance, dummy for take-up of commuter tax allowance, and dummy for take-up above the threshold 25,855 DKK

	Distance (km)		Commuter tax allowance (DKK)		Dummy (take-up vs no take-up)		Dummy (take-up more than 25,855 vs less)	
Reference average (2001-2003)	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
Treatment average (2004-2006)	0.160**	0.067	114.12**	41.64	-0.006**	0.002	0.003**	0.001
Treatment 2004	0.091	0.071	33.67	31.94	-0.007**	0.002	0.001	0.001
Treatment 2005	0.227**	0.080	94.02**	45.96	-0.006**	0.002	0.003**	0.001
Treatment 2006	0.162	0.100	138.54**	46.69	-0.005**	0.003	0.005**	0.001
N treatment: 82,590 control: 42,272 weighted individuals								

Notes: ** indicates that the coefficient is significant at the 5 % level, and * indicates significance at the 10 % level. Reference average (2001-2003) refers to the average of the three years before the reform. Treatment average (2004-2006) is the average of three years after the reform, whereas treatment in a given year compares the reference average2001-2003 to the yearly means.

5.2 The average treatment effect on the treated

The reform aimed at increasing workers' incentive to search for jobs farther away from home and thus expand the labor market to a larger area. As a consequence of the larger labor market, unemployment and poverty could decrease in the treated areas and not only benefit adults but as research has shown also the children (Chetty et al., 2015). We have already seen that the reform did indeed increase distance to work, however only for a small group. The question is if the impact of the reform was large enough to also drive down unemployment.

Table 7 presents the estimates of the effect of the commuter tax allowance reform on unemployment and wages. In the main results in the first row of Table 7, I compare averages of the three years pre-reform to averages of three years post reform. In the remaining rows I also present effect sizes by year post-reform, still compared to the three years averages pre-reform.

In the first column, the estimates show that the reform shows no indication of decreasing unemployment. Contrary the average estimate is insignificant and positive. As Table 7 also shows, the unemployment estimate is only negative in the first year after the reform, in which the insignificant estimate is -0.87.

Theoretically the reform could allow employers to lower their wage offer, but taken the size of the reform into account with an annual income increase of 0.5 percent, this is an unlikely scenario. The column for log (hourly wages) in Table 7 also shows an insignificant estimate of -0.32, i.e., from before to after the reform the wages decreased with 0.32 percent (in 2006 0.47 percent).

In conclusion, the extra commuter tax allowance increases distance to work with about 2.9 percent, 10.3 percent increase in residents reporting commuter tax allowance above the threshold of 50 km., but the allowance has not been enough to descrease unemployment. The lack of impact could be explained by either the small increase in the allowance (around 0.5 percent extra income for commutes above 50 km) or the high threshold of 50 km to work (for commuters outside their own municipality the average distance is around 26 km). As the subsidy is only transferred to residents working, the extra allowance might maintain some working residents from moving (even though the relocation has not changed, see Appendix B).

Even if that is the case, politicians must consider whether other types of local subsidies might be more cost-efficient in preventing depopulation of outlying municipalities, such as local tax deductions for all working residents, direct transfers to local businesses, or decentralizing governmental work places. Furthermore the national target to maintain a high degree of mobility among workers may be better met with other incentives, such as better public transportation or more public housing close to the largest labor market. The nationally and regionally common target of including more residents in the labor market has been unsuccessful through the increased commuter tax allowance. Whether the residents within-the-employed increase in distance to work is large enough to be cost-effective still remains.

Disregarding the the costs of implementing the reform, from 2003 to 2006 the extra spending among the treated per year was DKK 15.6 mio. For all residents in the treatment group, a total of 82,590 residents, the average extra income amounted to DKK 189 in 2004 prices (about 34 dollars in 2015). This back-of-the-envelope estimate for the costs of the reform shows that the an increase of 0.160 km, which is a 2.9 percent increase in average distance to work, costs DKK 189 per resident. This finding suggest that residents are responding quite a lot for a relatively small at a relatively small public costs. About 25 percent of residents take up the allowance, and the allowance has not been able to decrease unemployment. With this in mind, this extra spending are favoring an already employed population, and however small the costs, at least the regional target may not be met.

Table 7: Average treatment effect on the treated: unemployment and log(hourly wages)

	Unemployment		Log(hourly wages)	
Reference average (2001-2003)	Coef.	SE	Coef.	SE
Treatment average (2004-2006)	0.700	0.50	-0.023	0.02
Treatment 2004	-0.087	0.69	-0.012	0.02
Treatment 2005	1.223	1.76	-0.009	0.03
Treatment 2006	0.964	0.75	-0.047	0.03
N treatment: 82,590 control: 42,272 weighted individuals				

Notes: ** indicates that the coefficient is significant at the 5 % level, and * indicates significance at the 10 % level. Reference average (2001-2003) refers to the average of the three years before the reform. Treatment average (2004-2006) is the average of three years after the reform, whereas treatment in a given year compares the reference average2001-2003 to the yearly means.

6. Heterogenous effects

6.1 Distance to work

Despite the reform having limited effect on unemployment in outlying municipalities, it did significantly increase the distance to work, however the effect is small, implying that the job search area increased for some residents. Further analyzing who the reform affects, a dummy variable divides the sample into different types of commuters, assuming the value 1 for different distances to work in 2003, and zero otherwise. Commuters willing to commute before the reform might also be more willing to commute even farther as a consequence of the reform.

Table 8 presents the ATET estimates for the different subgroups of commuting distances before the reform: not commuting or commuting within own municipality (which is not detected in the distance variable), distance of 1-12 km, 13-20 km, 21-30 km, 31-40 km, 41-50 km, 51-60 km, and distance of more than 60 km.

The first row in Table 8 compares the average commuting distances three years before the reform to three years after the reform. Four significant estimates emerge. The first is for residents not commuting or commuting within own municipality, who increased their distance to work by 0.190 km after the reform. Compared to the average distance of 5.65 km in 2003, this is a 3.4 percent increase, and larger than the average effect of 0.160. As the estimates in the three post years separately reveals, the extra commuting for this group sets in in year 2005 and increase over time. The second estimate for residents commuting 13-20 km shows the same pattern, with a smaller estimate though. The distance for this group increased with 0.5 percent.

The third significant estimate is for residents commuting 31-40 km before the reform, just below the limit of 50 km. The negative estimate shows that these commuters reduced their commuting after the reform, apart from the first year after the reform. The fourth significant estimate exists for residents commuting more than 51-60 km, just above the limit of 50 km. These commuters, who have shown to have a taste for commuting, have decreased their distance to work after the reform. The effect is driven by year 2005 and 2006.

Table 8 demonstrates that the increase in distance to work is a result of increases at the intensive margin, and potentially also for residents driving longer than 60 km (positive but insignificant): residents who did not commute before the reform or only within own municipality, or relatively short distances between 21-30 km, and commuter who already commuted farther than 61 km. Commuters on the margin for the reform (51-60 km) decrease their distance to work. Even though Table 6 showed that fewer residents take up allowance after the reform, the reform has increased the distance to work for residents, who did not receive allowance before the reform. This could be explained by the renewed focus on commuter tax allowance, encouraging residents to find jobs father away from home. The distance estimate, however, reveal the same results at the self-reported commuter tax allowance showed: That the reform did increase commuting, but as the overall results also show, the increase was not strong enough to affect unemployment.

Table 8: Average treatment effect on the treated of commuters commuting 0, 1-12, 13-20, 21-30, 31-40, 41-50, 51-60 and more than 60 km before the reform on distance after the reform.

2003 distance	0 km (or own municipality)		1-12 km		13-20 km		21-30 km	
Reference average (2001-2003)	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
Treatment average (2004-2006)	0.190**	0.05	-0.001	0.01	0.031**	0.01	0.026	0.02
Treatment 2004	0.087	0.06	-0.019	0.01	0.004	0.02	-0.018	0.02
Treatment 2005	0.218**	0.07	0.002	0.02	0.046**	0.02	0.023	0.03
Treatment 2006	0.264**	0.07	-0.015	0.02	0.042**	0.02	0.071**	0.03
2003 distance	**31-40 km**		**41-50 km**		**51-60 km**		**≥ 61 km**	
Reference average (2001-2003)	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
Treatment average (2004-2006)	-0.048**	0.02	-0.003	0.02	-0.072**	0.01	0.038	0.03
Treatment 2004	0.004	0.02	0.036	0.02	-0.015	0.01	0.011	0.02
Treatment 2005	-0.053**	0.02	-0.001	0.02	-0.070**	0.01	0.062**	0.03
Treatment 2006	-0.094**	0.02	-0.044**	0.02	-0.131**	0.02	0.041	0.03

N treatment: 82,590 control: 42,272 weighted individuals

Notes: ** indicates that the coefficient is significant at the 5 % level, and * indicates significance at the 10 % level. Reference average (2001-2003) refers to the average of the three years before the reform. Treatment average (2004-2006) is the average of three years after the reform, whereas treatment in a given year compares the reference average2001-2003 to the yearly means.

6.2 Tied stayers—women in the outlying municipalities

As most individuals are part of a more than a one-person family, the decision of where to live and work becomes more complex. For couples, the search for jobs is geographically constrained, because jobs for both spouses must be within the same local labor market i.e., one spouse must compromise his or her job search. One type of model treats the household as one maximizing unit. When the couples decide on geographical location and jobs, they maximize their joint utility, choosing the location and job that amount to the highest joint utility.

Mincer (1978) finds that family ties often produce negative externalities in family member's job opportunities, mostly the wives', who experience a reduction in employment and in earnings. He calls these overqualified women "tied movers" (moving due to their family ties) and "tied stayers" (staying due to family ties). As Mincer (1978) and others show, the joint utility model has long been disputed by the empirical literature. Frank (1978) proposes a different approach. According to the "male chauvinistic family decision rule," the wife accompanies her husband to the geographical location that maximizes his job search, and from the given residential location she maximizes her own job search (Frank, 1978). The model explains the differential over-qualification of women in outlying areas.

For women already located in the treated outlying municipalities, the extra commuter tax allowance could potentially compensate them for longer commuting distances and expand their labor markets. However, if women are tied to home due to extra household responsibilities and child rearing, the incentive to drive farther to work may not be an option, and women will not respond to the reform. Black, Kolesnikova and Taylor (2014) find that the labor force participation rates of married women are negatively correlated with commuting time in metropolitan areas. Compton and Pollak (2014) support this finding. Estimating the effect of proximity to mothers or mothers-in-law on the labor supply of women with children, they find that women with children increase their labor supply if they have nearby family support for childcare.

If the reform had been large enough to attract long- distance commuting men and their spouses to the treated areas the reform could provide a negative effect on women's possibilities on the job market (tied movers). This would have been an unintended negative externality of the reform, and one policy makers should take into account, if they consider enlarging the reform in the future.

In the design of the empirical analyses, I can only examine tied stayers, not tied movers, since the treatment is being based on a treated municipality in 2003, and if one relocates to the municipality later, despite that this is a consequence of the reform, they are hidden in the control group. As women are often tied spouses, especially if they have children, I explore the heterogeneous effects separately for men and women at different ages, with and without children, and if they have children, whether they have a mother living in the same municipality to see if the reform has expanded the labor market for tied stayers.

There are many explanations why women are tied to the home. Madden and Chiu (1990) establish that women, and especially married women, have shorter commutes, and speculate that this is due to a relatively higher degree of household responsibilities than their husbands, lower potential earnings (which could be significant in Denmark, a country with a highly gender segregated labor market), and restricted vehicle access, which would also be more pronounced in the outlying municipalities, as most Danish families have only one car (60 percent of Danish families own at least one car, (Statistics Denmark, 2013). The extended household responsibilities of women have been established in several time use surveys (see an excellent study by Hochchild and Machung, 2012). If extended household responsibilities keep women from searching for jobs far from home, the extra commuter tax allowance could potentially compensate for the extra time used on commuting, but only if others can take over the responsibilities.

Grandparents are one widely used extra help in the home. In the Denmark, about 65 percent of grandmothers provide some sort of childcare for their grandchildren. Of the 65 percent, only about 20 percent provide childcare weekly or more often (Jappens and Bavel, 2012).

Table 9: Subsample analyses on women's distance to work and unemployment: overall, younger and older, with and without children, with children and mother in same municipality, and with children and without mother in same municipality.

	Women's unemployment		Younger, born 1964-1974		Older, born 1948-1963		With children	
Reference average (2001-2003)	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
Treatment average (2004-2006)	0.20	0.38	-0.13	0.30	0.46	0.32	0.27	0.39
Treatment 2004	-0.18	0.59	-0.28	0.41	0.11	0.40	-0.16	0.49
Treatment 2005	0.76	0.64	0.14	0.41	0.62	0.47	0.74	0.54
Treatment 2006	0.41	0.59	-0.24	0.40	0.65	0.42	0.22	0.51

	Without children		With children and mother nearby		With children and mother distant		Distance to work	
Reference average (2001-2003)	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
Treatment average (2004-2006)	0.06	0.24	-0.04	0.17	0.28	0.30	0.08*	0.04
Treatment 2004	-0.02	0.34	-0.21	0.21	0.05	0.42	0.03	0.05
Treatment 2005	0.02	0.33	0.23	0.23	0.50	0.46	0.12**	0.05
Treatment 2006	0.19	0.29	-0.14	0.25	0.36	0.45	0.08	0.06

Notes: ** indicates that the coefficient is significant at the 5 % level, and * indicates significance at the 10 % level. Reference average (2001-2003) refers to the average of the three years before the reform. Treatment average (2004-2006) is the average of three years after the reform, whereas treatment in a given year compares the reference average2001-2003 to the yearly means.

Table 9 shows these heteregenous effects. The most obvious measure in the table is the only significant one – distance to work. Men and women seem to have equally increased their distance to work (see appendix E, table E.1, for men). The remaining estimates are insignificant, and thus the signs could potentially be random. With this in mind, I cautiously interpret signs of the average estimates, as they seem to have the expected signs. Table 9 supports that younger women are more mobile than older. This distinction is not present among men. The sign is also negative for mothers with children, if their mothers live in the same municipality supporting that the reform only works, if the mothers have acces to after hours care for their children. The sign is large and positive and positive, for women with children with a mother in a different municipality. Even though the estimates are insignificant, they support the theoretical predictions.

7. Conclusion

Commuter tax allowance subsidizing commutes to and from work is a widely used policy to fulfill a national target that aims at keeping the labor force as mobile as possible, and through this mobility ensuring filling vacancies with the right person for the right job. Even though the policy is costly (0.2 percent of the Danish GDP in 2002; Statistics Denmark), little is known about the effects of the subsidy. Easing transportation barriers to jobs also benefit the targets of outlying localities. The local target is to attract and maintain a working tax-paying population, but with easing transportation barriers, the subsidy meet both targets – locally especially if the subsidy is place-based and thus is only given to residents in the given locality. This article adds to the literature by estimating the effect of an increase in commuter tax allowance targeted outlying municipalities. Using a reform in 2004 and matching estimatiors for identification in this natural experiment, I find the average treatment effect on the treated on distance to work, unemployment and wages.

I show that the introduction of the extra allowance increased the distance to work by 0.16 km, a 2.9 percent increase. The increased commuter tax allowance has, however, not been enough to decrease unemployment. The take-up of commuter tax allowance increased more than the distance to work, suggesting that the residents knew about the reform, but also reported more allowance the more they became aware of the possibility. This rather small effect should be seen in light of size of the reform.

The effect of the increased allowance depends on both the fiscal size of the commuter tax allowance and the threshold for obtaining the allowance. For the former, the amount of money received by residents must be worth at least the value of the time they spend on additional commuting. For residents commuting 60 km to and from job, the reform increased their annual income with about 0.6 percent. For the latter, the extra tax allowance has to be within reach, that is, the distance they drive to work in order to obtain the allowance has to fit into their lives in a viable way. To obtain the extra allowance the residents would have to commute longer than 50 km to and from work. Either the fiscal returns are too low or the threshold to obtain the allowance is too high, which could explain the small effect the reform has had.

The estimating method of exact propensity score matching demands that all relevant confounders of living in either the treated municipality or an adjacent municipality must be part of the propensity score. Failing to match on wages the years before the reform – though showing that the difference between the treated residents and their matched compairsons has been stable over the years before the reform – more research has to be done to conclude that an extra commuter tax allowance increase distance to work.

Performing subsample analyses searching for heterogeneous effects, the results are inconclusive. Deconstructing the distance to work effects by categories of former commuting, the results show that the residents driving less than 30 km were the ones increasing the distance after the reform. This suggests that the reform did not only increase commutes above the threshold of 50 km, but also informed the residents about commuter tax allowance in general. The largest increase in commuting was among residents reporting commuter tax allowance above the 50 km threshold. Exploring women's unemployment separately from men's, the effect on distance to work is significant, and the same for men and women. I find no significant differences on the remaining outcomes in the subsample analyses.

Acknowledgements

My thanks to my supervisors Marianne Simonsen and Nabanita Datta Gupta, and Rune Majlund Veilin, Rune Vammen Lesner, and Michael Svarer for valuable inputs and comments, and to the AU labour and policy seminar, SFI Advisory board seminar, especially Dean Lillard, and my co-supervisor at SFI, Beatrice Schindler Rangvid. I also thank Mona Larsen, SFI, and Jack Daniel Groth-Hinnerup for support. I gratefully thank SFI and AU RECEIV center (project no. 908792) for financial support.

Appendix A. An introduction to the Danish commuter tax allowance system

Contrary to many other countries, the Danish commuter tax allowance is self-reported on the annual tax returns. In the US the employer reports the commuter tax allowance and is also responsible for at least a percentage of the payment to the employee.

In Denmark, the individual reports his or her commuter tax allowance as often as he or she prefers and pays taxes accordingly. Once a year, at the annual tax return, the individual is responsible for reporting the annual tax allowance. Then, automatically, the tax returns are calculated, and the individual receives the excess tax payments back or pays the owed taxes. The tax returns are paid by the government.

The commuter tax allowance depends on the distance to work and only for commuters commuting farther than 12 kilometers to work, calculated in traveled kilometers and not as a straight line between home and work. Calculating the distance, the mode of transportation is irrelevant. The individual is only allowed to receive allowance on the actual commutes. If the IRS suspects an individual of tax fraud related to the commuter tax allowance, the individual has to show receipts and justify every trip. The burden of proof rests with the individual.

Commuters traveling by ferry, planes, crossing bridges with road taxes, as well as traveling abroad and staying overnight can receive extra deductions. Again, the individual who reports the commuter tax allowance is responsible for saving receipts, etc., for five years in case the IRS wants validation of the commute.

The employer can pay for employee transportation, but then the employee cannot receive commuter tax allowance for the commutes.

Appendix B. Relocation as an outcome

In this section, I analyze the effect of the increased commuter tax allowance on relocation. The reason for doing so is that if residents from adjacent municipalities moved to the treated municipalities due to the reform, in the design of the paper, this could potentially disguise an effect on unemployment. For example, if an employed control residents finds a job, because the resident now has an incentive to find a job farther away after moving, this would be a direct and intended effect of the reform. However, as this resident would be in the control group, this would diminish the effect of the reform on unemployment. For this to be a potential disguise, however, the residents in adjacent municipalities have to relocate more often than residents in the treatment. Therefore, I examine the effect of the reform on relocation.

The registers hold information on relocations of individuals from one address to another from one year to another. After matching, the residents in the treated and in the control groups have similar moving patterns.

To meet the requirement with similar trends in the treatment and control groups, Figure B.1 illustrates the differences in the yearly means between the treated and their matched comparisons in relocations. The trends seem to be similar in the two groups back to 1998. The graph flucturates around zero, which indicates that there are no signifiacant differences between the two groups. Figure B.1 supports that the matched difference-in-difference results have a causal interpretation.

Figure B.1: Relocation, before and after the reform in the treatment and control groups, respectively

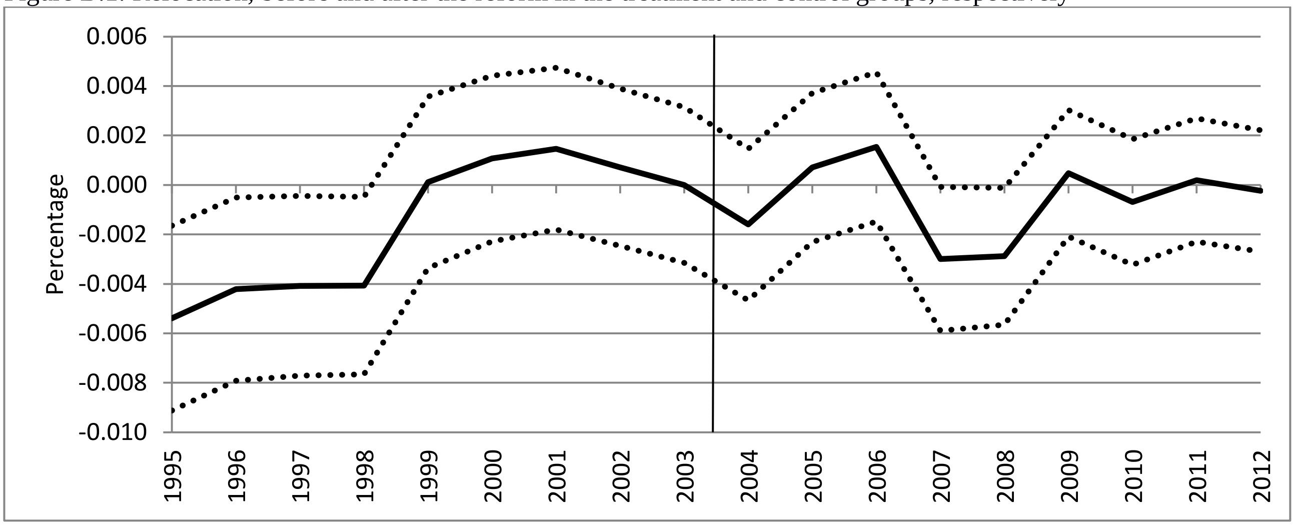

Table B.2 shows the estimation results, and I find no significant differences in relocation patterns from before to after the reform. From these findings, t he conclusion is that the lack of effect on unemployment cannot be due to residents from adjacent municipalities moving to the treated municipalities.

Table B.2: Average treatment effect on the treated: Relocation

	Relocation	
Reference average (2001-2003)	Coef.	SE
Treatment average (2004-2006)	-0.0001	0.002
Treatment 2004	-0.0022	0.003
Treatment 2005	0.0006	0.003
Treatment 2006	0.0020	0.003
N treatment: 82,590 control: 42,272 weighted individuals		

Notes: ** indicates that the coefficient is significant at the 5 % level, and * indicates significance at the 10 % level. Reference average (2001-2003) shows that comparison is the average of the three years before the reform. Treatment average (2004-2006) is the average of three years after the reform, whereas treatmentYEAR compares the before averages to the outcome in the given year.

Appendix C. Probit estimates used in the propensity score matching

Table C.1: Probit estimates used in the propensity score matching

Covariates		Coef.	SE
Commuter tax allowance			
	2003	0.005**	0.0005
Dummy for commuter tax allowance			
	2000	0.025**	0.009
	2001	0.043**	0.010
	2003	0.075**	0.009
Distance to work (km)			
	2000	0.0005**	0.0002
	2001	-0.0006**	0.0003
	2002	-0.0003**	0.0003
	2003	0.0013**	0.0003
Unemployment			
	1995	0.00005**	0.00002
	1996	0.00010**	0.00002
	1997	0.00004¤	0.00003
	1998	0.00005¤	0.00003
	1999	0.00009**	0.00003
	2000	0.00016**	0.00003
	2001	0.00014**	0.00003
	2003	0.00012**	0.00003
Dummy for relocation			
	2002	-0.011**	0.008
	2003	-0.036*	0.008
Log (annual wages)			
	2000	-0.0007**	0.001
	2001	-0.0046¤	0.001
Woman		-0.027**	0.006
Other ethnicity than Danish		-0.057**	0.013
Family type 2003			
	Married (left out)		
	Registered partnership	-0.153**	0.095
	Cohabiting with children	0.027**	0.010
	Cohabiting without children	0.004	0.009
	Single	0.039**	0.007

Table C.1 *(continued)*: Probit estimates used in the propensity score matching

Covariates	Coef.	SE
Education in 2003		
No education (left out)		
High school	0.031**	0.006
Vocational	-0.010	0.012
Short or medium	-0.022	0.010
Bachelor	-0.129**	0.034
Master or PhD	-0.162**	0.016
Education Missing	-0.039	0.023
Sector at work 2003		
Missing sector (left out)		
Integrated governmental	-0.024*	0.020
Integrated non-governmental	-0.106	0.069
Non-integrated quasi governmental	0.090**	0.021
Social foundations	0.094**	0.158
Integrated regional/municipal	0.020**	0.153
Integrated non-regional	-0.496**	0.075
Non-integrated quasi regional	-0.177*	0.051
Integrated municipal	0.039**	0.012
Integrated non-municipal	-0.117*	0.035
Non-integrated quasi municipal	-0.002	0.035
Integrated governmental	-0.175	0.066
Integrated regional	0.028¤	0.195
Integrated primarily municipal	-0.202*	0.062
Non-integrated public	-0.134	0.026
Private	-0.047*	0.010
Employment status		
No employment (left out)		
Weekly work hours>27, insured1	-0.018**	0.011
Weekly work hours>27, uninsured	-0.055¤	0.016
Weekly work hours>27, insured2	0.200	0.319
Part-time insured	0.205**	0.032
Weekly work hours 18-27 hours, uninsured	0.011**	0.028
Weekly work hours 9-18 hours, uninsured	-0.103	0.036
Weekly work hours 1-9 hours, uninsured	-0.072	0.038
Weekly work hours unknown, uninsured	-0.017	0.032
Missing information	-0.053	0.050

Table C.1 *(continued)*: Probit estimates used in the propensity score matching.

Covariate	Coef.	SE
Dummy for cohorts		
1951-1955	-0.005	0.008
1961-1965	-0.015*	0.008
1966-1970	-0.030**	0.008
1971-1976	-0.035**	0.008
Region		
Northern Jutland (left out)		
Middle Jutland	-1.076**	0.001
Southern Jutland	-0.933**	0.001
Fuen	1.097**	0.011
Constant	-0.043 **	0.015
Observations		284,577

Notes: ** indicates that the coefficient is significant at the 5 % level, * significance at the 10 % level, and ¤ significance at the 20 % level

Appendix D. Kernel distrubition of propensity scores and balancing tests

Figure D.1: Kernel distribution of propensity scores for Northern Jutland before and after propensity score matching

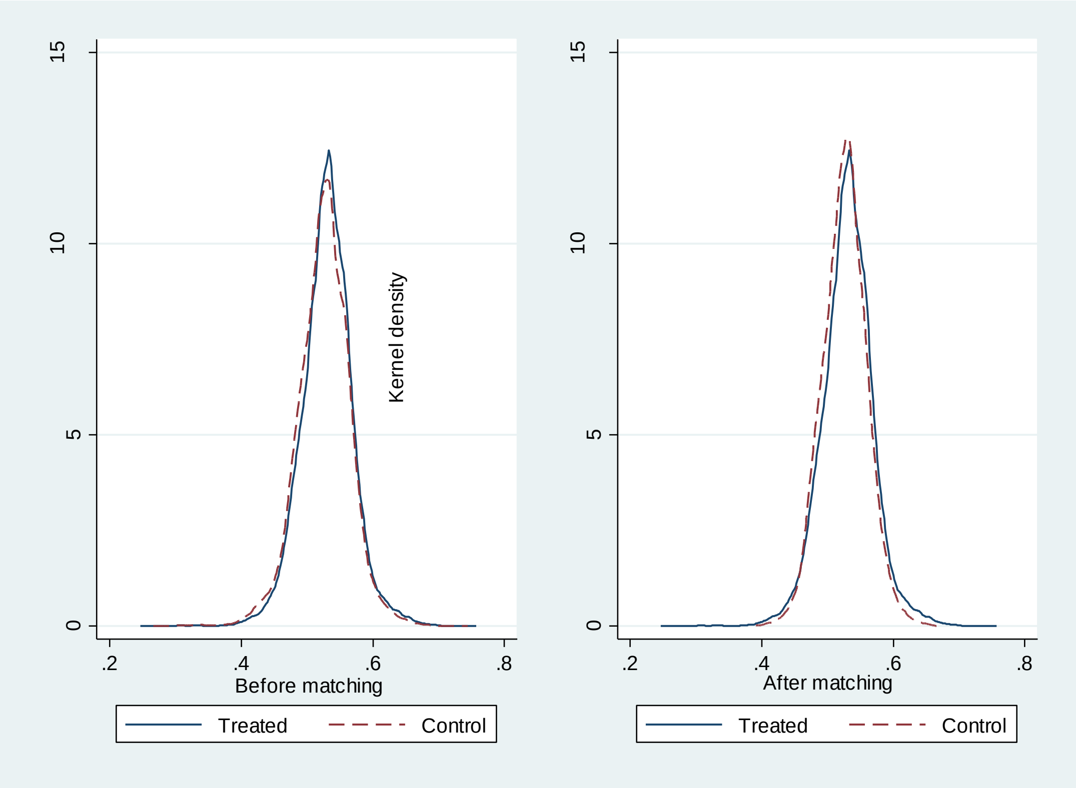

Figure D.2: Kernel distribution of propensity scores for Middle Jutland before and after propensity score matching

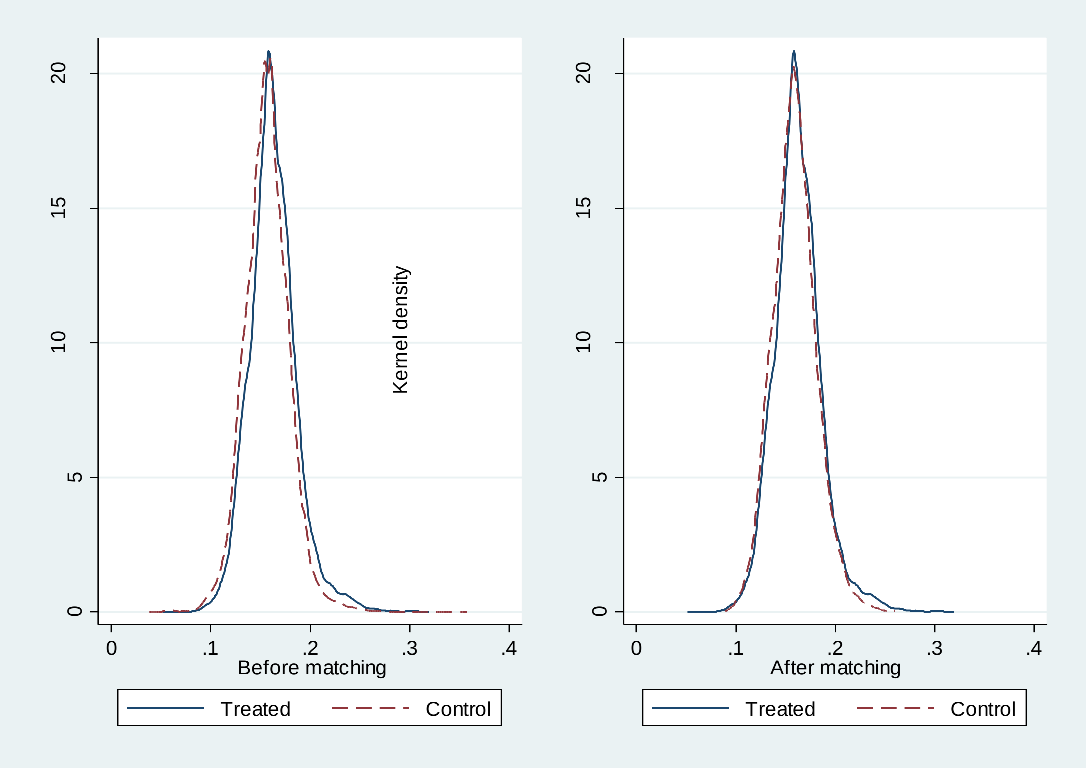

Figure D.3: Kernel distribution of propensity scores for Southern Jutland before and after propensity score matching

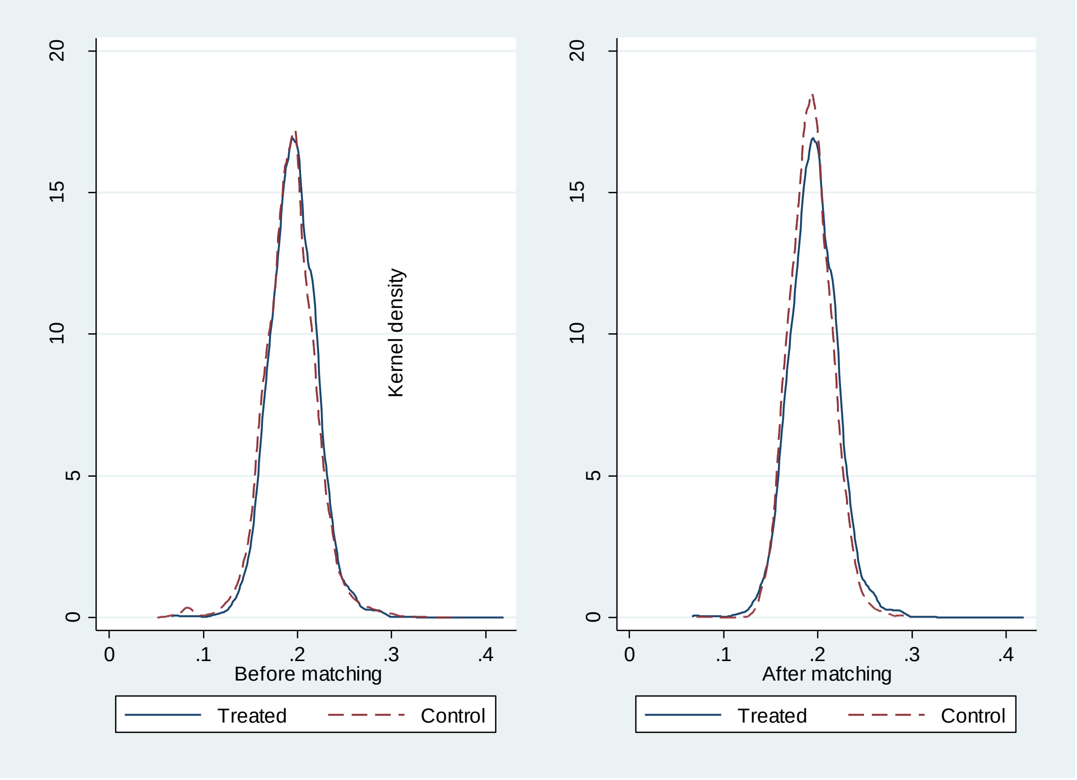

Figure D.4: Kernel distribution of propensity scores for Fuen before and after propensity score matching

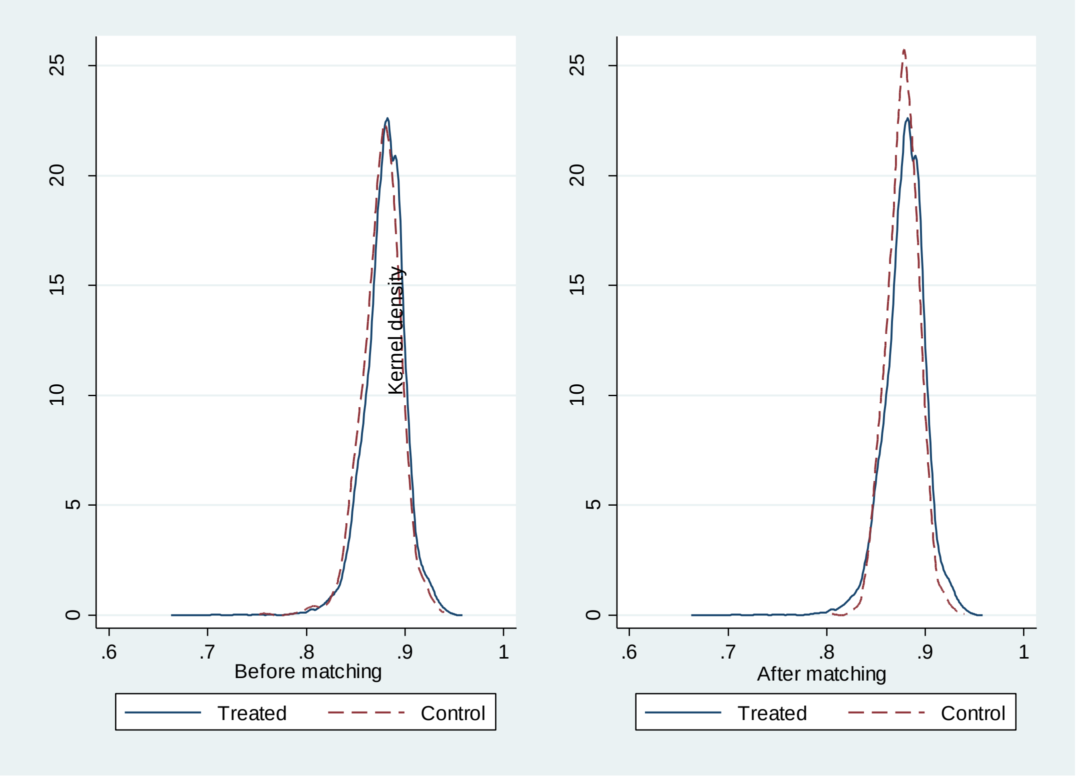

Table D.1. Balancing tests: Means of the treatment and control group, and a t-test to compare the means

Variable	Year	Before or after matching	Mean Treated	Mean Control	%bias	% bias reduction	t- t-value	
Distance to work (km)	2000	Before	5.47	5.47	0.00		0.05	
		After	5.12	5.18	-0.40	-1973	-0.91	
	2001	Before	5.64	5.65	-0.10		-0.25	
		After	5.25	5.37	-0.80	-676	-1.73	
	2002	Before	5.68	5.68	0.00		-0.10	
		After	5.35	5.44	-0.60	-1495	-1.42	
	2003	Before	5.91	5.70	1.40		3.71	**
		After	5.46	5.49	-0.20	86	-0.47	
Unemployment	1995	Before	35.04	29.34	4.80		12.46	**
		After	31.41	32.11	-0.60	88	-1.21	
	1996	Before	35.71	29.07	5.70		14.87	**
		After	30.29	30.35	-0.10	99	-0.12	
	1997	Before	30.96	26.06	4.50		11.71	**
		After	26.98	27.12	-0.10	97	-0.28	
	1998	Before	28.57	23.76	4.70		12.29	**
		After	24.70	24.41	0.30	94	0.61	
	1999	Before	28.30	22.95	5.40		14.13	**
		After	23.43	23.50	-0.10	99	-0.15	
	2000	Before	29.85	23.33	6.30		16.40	**
		After	23.04	23.60	-0.50	92	-1.22	
	2001	Before	29.50	22.90	6.50		17.09	**
		After	23.00	23.49	-0.50	93	-1.08	
	2002	Before	25.15	20.92	4.50		11.68	**
		After	22.46	22.68	-0.20	95	-0.48	
	2003	Before	27.47	22.98	4.40		11.42	**
		After	23.33	24.12	-0.80	82	-1.66	
Log (annual wages)	2000	Before	10.10	10.33	-5.50		-14.04	
		After	10.06	9.99	1.90	66	**3.63**	**
	2001	Before	10.17	10.40	-5.40		-13.71	**
		After	10.12	10.04	2.00	64	**3.82**	**
	2002	Before	10.18	10.42	-5.60		-14.25	**
		After	10.13	10.06	1.80	67	**3.60**	**
	2003	Before	10.16	10.41	-6.00		-15.40	**
		After	10.11	10.05	1.50	75	**2.97**	**

Table D.1 *(continued)*: Balancing tests: Means of the treatment and control group, and a t-test to compare the means

Variable		Year	Before or after matching	Mean Treated	Mean Control	%bias	% bias reduction	t-value	
Dummy for relocation		2000	Before	0.14	0.15	-0.90		-2.18	**
			After	0.14	0.14	0.20	82	0.33	
		2001	Before	0.14	0.14	-0.90		-2.41	**
			After	0.13	0.14	-1.00	-9	**-2.12**	**
		2002	Before	0.13	0.13	-0.10		-0.30	
			After	0.12	0.12	0.20	-66	0.40	
		2003	Before	0.13	0.13	-0.90		-2.25	**
			After	0.12	0.12	-0.10	87	-0.24	
Woman			Before	0.49	0.49	0.40		0.91	
			After	0.49	0.49	0.50	-38	1.00	
Commuter tax allowance (DKK 2004 prices)		2000	Before	3093.20	3274.00	-2.60		-6.46	**
			After	3075.00	3041.50	0.50	81	0.97	
		2001	Before	3257.40	3446.90	-2.60		-6.54	**
			After	3231.20	3193.70	0.50	80	1.05	
		2002	Before	3452.50	3619.60	-2.20		-5.59	**
			After	3418.80	3450.90	-0.40	81	-0.86	
		2003	Before	3601.50	3727.80	-1.60		-4.18	**
			After	3541.20	3589.70	-0.60	62	-1.30	
Dummy for commuter tax allowance		2000	Before	0.71	0.69	4.30		10.98	**
			After	0.72	0.72	0.00	100	-0.04	
		2001	Before	0.71	0.69	4.60		11.74	**
			After	0.71	0.71	0.10	98	0.18	
		2002	Before	0.71	0.68	4.70		12.06	**
			After	0.71	0.71	0.90	81	1.84	
		2003	Before	0.69	0.67	4.20		10.77	**
			After	0.69	0.69	0.70	84	1.43	
Education	No education (left out)								
	High school		Before	0.44	0.43	1.90		4.85	**
			After	0.45	0.45	-0.10	97	-0.13	
	Vocational		Before	0.06	0.06	-1.00		-2.45	**
			After	0.06	0.06	-0.30	71	-0.58	
	Short or medium		Before	0.13	0.13	-0.50		-1.40	
			After	0.13	0.12	2.00	-263	**4.09**	**
	Bachelor		Before	0.01	0.01	-2.00		-5.12	**
			After	0.00	0.00	0.10	93	0.35	
	Master or Phd		Before	0.03	0.04	-4.00		-10.09	**
			After	0.02	0.02	-1.00	76	**-2.53**	**
	Education missing		Before	0.01	0.01	-0.40		-1.11	
			After	0.01	0.01	0.40	9	0.83	

Table D.1 *(continued)*: Balancing tests: Means of the treatment and control group, and a t-test to compare the means

Variable	Year	Before or after matching	Mean Treated	Mean Control	%bias	% bias reduction	t-value	
Other ethnicity than Danish		Before	0.04	0.05	-3.90		-9.88	**
		After	0.04	0.05	-0.40	90	-0.82	
Dummy for children in 2003		Before	0.60	0.60	-1.10		-2.93	**
		After	0.60	0.60	-0.20	79	-0.48	
Family type in 2003	Married (left out)							
	Reg. partnership	Before	0.00	0.00	-0.40		-0.89	
		After	0.00	0.00	0.20	39	0.53	
	Cohabiting with children	Before	0.09	0.09	1.90		4.85	**
		After	0.09	0.10	-1.10	44	**-2.13**	**
	Cohabiting without children	Before	0.10	0.10	-0.10		-0.25	
		After	0.10	0.10	0.50	-392	0.99	
	Single	Before	0.23	0.22	2.20		5.61	**
		After	0.22	0.22	1.10	51	**2.19**	

Notes: ** indicates that the means differ, and bold that means differ in the matched sample. % bias is the standardized percentagebias after the matching, and a value above 1 suggest that the matching was unsuccessful. % bias reduction is the difference of the sample means in the treated and non-treated subsamples, as a percentage of the square root of the average of the sample variances in the treated and the non-treated groups (Stata-help, ptest).

Appendix E. Subsample analyses on men

Table E.1: Subsample analyses on men's distance to work and unemployment: overall, younger and older, with and without children, with children and mother in same municipality, and with children and without mother in same municipality.

	Overall		Younger, born 1964-1974		Older, born 1948-1963		With children	
Reference average (2001-2003)	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
Treatment average (2004-2006)	0.51	0.34	0.29	0.22	0.22	0.24	0.25	0.27
Treatment 2004	0.45	0.54	0.33	0.35	0.12	0.37	0.29	0.38
Treatment 2005	0.51	0.42	0.29	0.27	0.22	0.33	0.07	0.30
Treatment 2006	0.58	0.41	0.25	0.31	0.33	0.31	0.39	0.32

	Without children		With children and mother		With children and mother distant		Distance to work	
Reference average (2001-2003)	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
Treatment average (2004-2006)	0.26	0.24	0.05	0.09	0.28	0.30	0.08	0.06
Treatment 2004	0.12	0.39	0.16	0.12	-0.05	0.30	0.09	0.06
Treatment 2005	0.51	0.35	-0.04	0.13	0.16	0.25	0.08	0.08
Treatment 2006	0.27	0.38	0.02	0.11	0.39	0.27	0.08	0.08

N treatment: 82,590 control: 42,272 weighted individuals

Notes: ** indicates that the coefficient is significant at 5 % level, and * indicates significance at the 10 % level. Reference average (2001-2003) refers to comparison being the average of the three years before the reform. Treatment average (2004-2006) is the average of three years after the reform, whereas treatmentYEAR compares the before averages to the outcome in the given year.

References

Abadie, Alberto, and Guido W. Imbens, 2006. Large Sample Properties of Matching Estimators for Average Treatment Effects, *Econometrica* 74(1), 235-267.

Abadie, Alberto, and Guido W. Imbens. 2008. On the Failure of the Bootstrap for Matching Estimators. *Econometrica* 76(6), 1537–57.

Abadie, Alberto, and Guido W. Imbens. 2011. Bias-corrected matching estimators for average treatment effects. *Journal of Business & Economic Statistics* 29(1), 1-11.

Angrist, Joshua D., and Jörn-Steffen Pischke, 2008. Mostly Harmless Econometrics: An Empiricists' Companion, Princeton University Press, Princeton, NJ.

Caliendo, Marco, and Sabine Kopeinig, 2008. Some practical guidance for the implementation of propensity score matching. *Journal of economic surveys* 22(1): 31-72.

Chetty, Raj, Nathaniel Hendren, and Lawrence F. Katz, 2015. The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment. *National Bureau of Economic Research*, Working paper 21156.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, 2004. How Much Should We Trust Differences-in-Differences Estimates? *The Quarterly Journal of Economics* 119 (1), 249-275.

Black, Dan A., Natalia Kolesnikova, and Lowell J. Taylor. Why do so few women work in New York (and so many in Minneapolis)? Labor supply of married women across US cities. *Journal of Urban Economics* 79, 59-71.

Blundell, Richard, and Monica Costa Dias, 2009. Alternative approaches to evaluation in empirical microeconomics. *Journal of Human Resources* 44(3), 565-640.

Borck, Rainald, and Matthias Wrede, 2005. Political economy of commuting subsidies. Journal of Urban Economics 57(3), 478-499.

Busso, Matias, and Patrick Kline, 2008. Do Local Economic Development Programs Work? Evidence from the Federal Empowerment Zone Program. *Cowles Foundation* Discussion Paper 1638. Cowles Foundation, Yale University.

Compton, Janice, and Robert A. Pollak, 2014. Family proximity, childcare, and women's labor force attachment. *Journal of Urban Economics* 79, 72-90.

Frank, Robert H, 1978. Why women earn less: the theory and estimation of differential overqualification. *The American Economic Review* 68(3), 360-373.

Givord, Pauline, Roland Rathelot, and Patrick Sillard, 2013. Place-based tax exemptions and displacement effects: An evaluation of the Zones Franches Urbaines program. *Regional Science and Urban Economics* 43(1), 151-163.

Gobillon, Laurent, and Harris Selod., 2014. Spatial mismatch, poverty, and vulnerable populations. *Handbook of Regional Science*, 93-107. Springer Berlin Heidelberg.

Harding, M. , 2014. Personal Tax Treatment of Company Cars and Commuting Expenses: Estimating the Fiscal and Environmental Costs. *OECD Taxation* Working Papers 20, OECD Publishing.

Haas, Anette, and Liv Osland, 2014. Commuting, migration, housing and labour markets: Complex interactions. *Urban Studies* 51(3), 463-476.

Heckman, James J., and Jeffrey A. Smith. 1999. The Pre-programme Earnings Dip and the Determinants of Participation in a Social Programme. Implications for Simple Programme Evaluation Strategies. *Economic Journal* 109(457): 313-48.

Heckman, James J., Hidehiko Ichimura, and Petra Todd, 1998. Matching as an econometric evaluation estimator. *The Review of Economic Studies* 65(2), 261-294.

Hochschild, Arlie, and Anne Machung, 2012. The second shift: Working families and the revolution at home. Penguin.

Ihlanfeldt, Keith R., and David. L. Sjoquist,1998. The spatial mismatch hypothesis: a review of recent studies and their implications for welfare reform. *Housing policy debate* 9(4), 849-892.

Imbens, Guido W., and Jeffrey M. Wooldridge, 2009. Recent Developments in the Econometrics of Program Evaluation, *Journal of Economic Literature, American Economic Association* 47(1), 5-86

Jappens, Maaike, and Jan Van Bavel, 2012. Regional family norms and child care by grandparents in Europe. *Demographic research* 27(4), 85-120.

Kain, John F. , 1968. Housing segregation, negro employment, and metropolitan decentralization. *The Quarterly Journal of Economics* 82(2), 175-197.

Katz, Lawrence F., Jeffrey R. Kling and Jeffrey B. Liebman, 2001. Moving To Opportunity In Boston: Early Results Of A Randomized Mobility Experiment, *Quarterly Journal of Economics* 116(2), 607-654

Kleven, Henrik Jacobsen, Martin B. Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez, 2011. Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark. *Econometrica* 79(3): 651-692.

Kleven, Henrik Jacobsen, Camille Landais, Emmanuel Saez, and Esben Schultz, 2014. Migration and Wage Effects of Taxing Top Earners: Evidence from the Foreigners' Tax Scheme in Denmark. *Quarterly Journal of Economics* 129: 333-378.

Lassen, David Dreyer, and Søren Serritzlew, 2011. Jurisdiction size and local democracy: evidence on internal political efficacy from large-scale municipal reform. *American Political Science Review* 105(2), 238-258.

Lechner, Michael, 2011. The Estimation of Causal Effects by Difference-in-Difference Methods. *Foundations and Trends in Econometrics* 4(3), 165-224.

Leth-Petersen, Søren, 2010. Intertemporal consumption and credit constraints: Does total expenditure respond to an exogenous shock to credit? *The American Economic Review* 100(3), 1080-1103.

Madden, J. F., & Chen Chiu, L., 1990. The wage effects of residential location and commuting constraints on employed married women, *Urban Studies*, 27, 353-369.

Manning, Alan; Petrongolo, Barbara , 2011. How local are labor markets? Evidence from a spatial job search model, IZA Discussion Paper No. 6178, Mincer, Jacob, 1978. Family migration Decisions. *Journal of Political Economy* 86(5), 749-773

Paetzold, Jürg, and Hannes Winner, 2014. Taking the High Road? Compliance with commuter tax allowances and the role of evasion spillovers. *Oxford University Centre for Business Taxation* Working papers 1419.

Patacchini, Eleonora, and Yves Zenou, 2005. Spatial mismatch, transport mode and search decisions in England. *Journal of Urban Economics* 58(1), 62-90.

Posadas, Josefina, and Marian Vidal-Fernández 2012. Grandparents' childcare and female labor force participation. *IZA* Discussion Paper No. 6398.

Roberts, Jennifer, Robert Hodgson, and Paul Dolan, 2011. “It's driving her mad: Gender differences in the effects of commuting on psychological health. *Journal of health economics* 30(5), 1064-1076.

Rosenbaum, James E., and Anita Zuberi, 2010. Comparing residential mobility programs: design elements, neighborhood placements, and outcomes in MTO and Gautreaux. *Housing Policy Debate* 20(1), 27-41.

Rouwendal, Jan, 2014. Commuting, Housing, and Labor Markets. *Handbook of Regional Science*, 75-91.

Russo, Giovanni, Federico Tedeschi, Aura Reggiani, and Peter Nijkamp, 2013. Commuter effects on local labour markets: a German modelling study. *Urban Studies* 51(3), 493-508.

Weinhardt, Felix, 2014. Social housing, neighborhood quality and student performance. *Journal of Urban Economics* 82, 12-31.

White, Michelle J., 1986. Sex differences in urban commuting patterns. The American economic review 76(2), 368-372.

Wrede, Matthias, 2001. Should commuting expenses be tax deductible? A welfare analysis. Journal of Urban Economics 49(1), 80-99.

Intra-household bargaining costs: The case of fathers' parental leave

Lisbeth Palmhøj Nielsen

Abstract This paper investigates the trajectories for two different systems of sharing parental leave: free-choice and father-quota. I discuss how and why intra-household bargaining over division of parental leave favors a larger maternal than paternal uptake of parental leave. Due to gendered spheres influencing bargaining, the mother receives a larger share by default, which in turn decreases the space for intra-household bargaining, and bargaining costs diminish this space even further. Gendered spheres bargaining, especially for high bargaining cost families, can explain why Danish men in a free-choice setting use only a small part of the entire parental leave available. It could also explain why fathers in countries with earmarked paternal leave take up more. Combining the Danish Longitudinal Survey of Children born in 1995 (DALSC) and Danish register data, the empirical findings support the hypothesis that high bargaining costs favor maternal take-up of parental leave at the expense of paternal take up.

Keywords Parental leave • Cost of bargaining • Intra-household • Bargaining power • Sharing rules

JEL Classifications D13 • J13

1. Introduction

As the only Scandinavian country, Denmark has decided against earmarking parental leave for the father, apart from the first two weeks after the birth of the child. Instead, Danish parents choose for themselves how to divide a percentage of the leave between them; the free-choice system. In 2011, Danish fathers took up 7.4 percent of the entire leave. In countries with so-called father quotas, the take-up was considerably higher. In 2011, fathers took up between 17.8 and 29 percent of the entire leave (Nordic Statistical Yearbook, 2012). The use-it-or-lose-it earmarked paternal leave appears to incentivize fathers to take more leave than they would have otherwise.

The literature shows women's, but not men's, career paths spiraling downward as soon as they have children (Drange and Rege, 2013; Waldfogel, 1998; Ejrnæs and Kunze, 2013). Even though women, in Scandinavia in particular, but also in other EU-countries and in the U.S., take up a large share of jobs on the labor market today, they still receive lower earnings than men (Waldfogel, 1998) and are responsible for the larger share of the household chores (Bianchi, Milkie, Sayer, and Robinson, 2000). Despite the generous Scandinavian family policies, mothers spend more time on household chores and child rearing than men, even after the period of maternal leave is over, and often reduce their working hours (Datta Gupta and Stratton, 2010).

Lundberg and Pollak (1993) introduce the gendered spheres bargaining model to explain the gendered differences in various chores and how couples in case of non-cooperation resort to their gendered chores. They also explain inefficiency in marriage being due to transaction costs from bargaining, which is here referred to simply as bargaining costs. Bargaining costs are emotional costs of negotiation between spouses, such as resentment, sulking, retaliation and in extreme cases domestic violence.

This paper discusses the differences between free-choice and quota systems and explores how gendered spheres and bargaining costs affect the division of parental leave. The theoretical model builds on the notion of gendered spheres, which means that the parental leave belongs to the mother a priori. For couples who disagree about how they prefer to share the leave, the gendered spheres bargaining setup leaves little room for negotiation. The couple can only negotiate parental leave if the mother is willing to give up a share of the leave and the father is willing to take up a share of leave. For couples with high bargaining costs, the gendered spheres result in even less leave for the father.

Empirically I test whether the presence of bargaining costs correlates with fathers' parental leave, and how relative bargaining power and the experience of the parents impact the results. Combined with Danish register data, I use the Danish Longitudinal Survey of Children – DALSC – that follows the 1995 birth cohort. Surveys in the literature rarely include fathers, but because I have access to 4003 couples where both parties have answered the questionnaires, the answers reveal information about the costs of bargaining for the whole family.

The contribution of this study is to guide policy makers to construct the most efficient parental leave system possible, that is, a system that produces Pareto efficient outcomes to the benefit of all family members.

2. Bargaining models

Becker (1981) was the first economist to introduce a model that could explain the behavior of households: the unitary model. In this model, the family has a joint utility function and the family resources are allocated by the altruistic head (or despotic dictator) of the family. Empirically, however, this type of Pareto efficient joint family utility function has been rejected (see among others Konrad and Lommerud 2000; Lundberg and Pollak 2003), because family members do not necessarily operate as one unit, but as individuals with private preferences – preferences that may coincide with other family members' utility.

2.1 Non-unitary models

The non-unitary models can be roughly divided into two intertwined categories: cooperative and non-cooperative models. The cooperative models have the attractive feature that they produce Pareto efficient outcomes, which is why many researchers have used them to explain household behavior (Browning and Chiappori, 1998), whereas the more complex non-cooperative models need not produce Pareto efficient outcomes.

In the cooperative models (see Chiappori (1988, 1992, 1997)), couples maximize a Nash production function of preferences for both spouses subject to a pooled budget constraint. Each spouse's preferences apply with different weights, which are determined by the relative bargaining power of the spouses, and the bargaining of the family is costless. If one spouse has relatively strong bargaining power, for instance by being better educated, earning higher income, appearing more attractive, etc., he or she may choose to disregard preferences of the other spouse and using the family solely for optimizing his or her own utility function. Within the frames of the marriage, the couple makes binding agreements. An important distinction from the unitary models is that each spouse has a lower limit as to how much abuse or inequality in consumption they are willing to tolerate: the threat point. The threat point in cooperative models is divorce, so if a spouse receives less utility than as single (the outside option), that spouse prefers divorce to staying in the relationship.

The non-cooperative models acknowledge the possibility that at least one of the spouses could behave inefficiently. The non-cooperative models differ from the Nash cooperative models, because they allow for inefficient outcomes due to bargaining itself being potentially costly (transaction costs in Lundberg and Pollak 1993). This may lead spouses to optimize individually, taking their spouse's inputs as given, and act subject to own income instead of pooled income. The non-cooperative models focus less on legally binding contracts and agreements, as the intra household cooperative models do, and rely instead on self-enforcing equilibria. In some cases, the cooperative and non-cooperative models will generate the same outcome, but usually not. Some models use individual optimization models, such as Cournot or Stackelberg in support of an inefficient allocation (Amilon, 2003), others use different threat points than divorce (Lundberg and Pollak, 1993). In the non-cooperative models, each spouse optimizes individually without taking the joint utility of family members into account, and therefore the solution is not Pareto efficient.

2.2 The non-cooperative gendered spheres model

To model how spouses decide on their respective contribution to parental leave in a free-choice system, I combine the base model from Amilon (2003) and the gendered spheres threat point by Lundberg and Pollak (1993).

In the base model of Amilon (2003), each spouse obtains utility from private consumption and the public good: the child. Each spouse decides how much time to invest in paid labor and parental leave respectively, but if they spend time on leave, they earn less income. Browning and Gørtz (2012) show that higher private income leads to higher private consumption. Consequently, if a spouse takes up a longer leave that spouse will lose private consumption. A strong preference for leave can compensate for the wage loss, but if the leave exceeds the preferred level the spouse on leave receives less utility from extra leave than on his or her paid job.

The costs of forgone private consumption may not be the only cost of being on leave. Additional costs may differ considerable for mothers and fathers. Close to all Danish women take up the fourteen weeks of maternity leave, and thus forego not only wages, but also the costs such as foregone visibility and good assignments on the job, and the employer suffer the cost of losing the employee for a considerable amount of time, having to find a substitute worker, rescheduling tasks etc. Taking a short leave of say two weeks may not cause any problem for the parent, but the longer the leave, the more costs up until some threshold. One would expect that exceeding the threshold, these marginal costs decline, and hereon after an extra week (or ten) is less costly. For the division of leave this is critical. The father may be willing to take the two earmarked weeks of father leave, but the costs following the two weeks on father leave are high. For the mother, who has already been on leave for 14 weeks, her marginal costs could be declining at this point. Even without gendered spheres, the one spouse on leave for fourteen weeks endures fewer costs by an extra week, and by construction the first 14 weeks are earmarked by the mother. As the costs are gendered, it points in the same direction as the gendered spheres: that the leave belongs to mother.

Even though non-cooperation may eventually lead to a divorce, divorce is not the most likely scenario right after having a baby, even if the spouses have conflicting interests when it comes to leave. Because the decision of parental leave-sharing only occurs a few times in life, divorce is unlikely based on this decision alone, even though the consequences of the leave may be severe for future income and bargaining power within the relationship.

In the case of non-cooperation, the couple is more likely to resort to gender-specific division of leave and gendered spheres bargaining as described in Lundberg and Pollak (1993). These gendered spheres affect the individual threat points, but differently for each gender and type of chore. Amilon (2007) uses gendered spheres to explain why temporary parental leave is usually taken by the mother. Here, I assume that gendered spheres exist within the spheres of parental leave, and that the parental leave by default belongs to the mother. In the case of non-cooperation between the spouses, the mother takes up all the leave and the father none.

Because the gendered spheres affect the threat point of each spouse, the gender of the spouse determines possibilities for bargaining. If the mother prefers to take all the leave, or the father prefers to take none, even though the spouses disagree, they cannot abstain from the default equilibrium that the mother takes up all the leave. With the gendered spheres threat point, the assumption is that non-cooperative couples still benefit from being married through joint consumption of the public good – in this case the public good is the child. One might think of these benefits as both parties consuming the public good all the time (both living with the child), instead of alternately in case of a divorce.

Lundberg and Pollak (1993) explain that the reason couples do not cooperate is the transaction costs of bargaining (bargaining costs). If a spouse knows that discussing an issue will cause a conflict, he or she might decide against it and stay within his or her own gendered domain. Still, if one spouse has relatively high bargaining power and expects to win the argument, the bargaining costs might be worthwhile. Bargaining power is still used in the non-cooperative models to determine default spheres and whether or not family members decide to cooperate in the first place. Lundberg and Pollak (1993) assume that either the couple uses a Nash cooperative bargaining process, or they refrain from cooperation, return to their gendered spheres, ending in a Cournot equilibrium, in which each spouse takes the contribution of the other as given, and then optimizes alone. The costs of bargaining amplify the impact of the gendered spheres threat point.

2.3 The free-choice system with gendered spheres

In a free-choice system, where spouses have to decide how to share parental leave, gendered spheres decreases the room for bargaining. The gendered spheres assign the leave to the mother by default, and thus her threat point can only be a minimum of leave she is willing to give up, T_{min}, but not a maximum of leave she is willing to take. For the father, the opposite argument applies: Since the leave does *not* belong to him by default, he can only set a threat point of the maximum of leave he is willing to take, T_{max}.

Figure 1 shows the room for bargaining in two different gendered spheres scenarios: at the top without bargaining costs, and at the bottom including bargaining costs. In both scenarios, the mothers' *potential* take-up of parental leave is shown on the y-axis and the fathers' on the x-axis. I use the Danish 1995 leave system as case, in which parents can take-up between zero and 10 weeks (see section 3.1). The full lines show the individual threat point of the mother, T_{min}, and the father, T_{max}. The black dots show the possible outcomes of leave, as I have assumed that the couple will use the full weeks of leave available for the benefit of the child. The question is how they are going to share the leave.

Figure 1: Gendered spheres bargaining, without bargaining costs

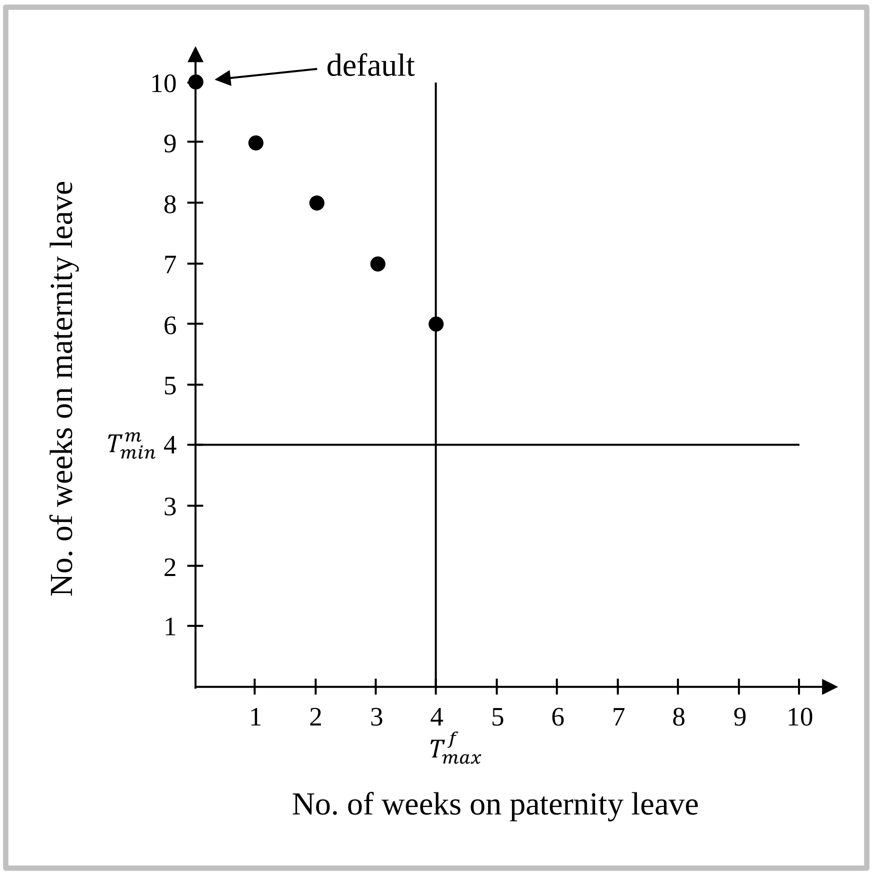

Even if the spouses have the same preference for leave and no bargaining costs, the gendered threat points push the window of opportunity towards the mother (and the opportunity for paid labor towards the father), as shown in the example in Figure 1. If they do not agree on a division within their joint threat points, they end up with a division of leave with all ten weeks for the mother and none for the father. This could potentially explain the low paternal take-up of parental leave in the free-choice regime of Denmark. Figure 1 shows how the differences in gendered threat points restrict the space for negotiation: If the mother sets her threat point T_{min} to ten weeks, the space for negotiation disappears, regardless of paternal preferences for leave. The opposite holds true for the father: If he sets his threat point T_{max} to zero weeks, there is no room for negotiation.

It becomes obvious that bargaining over parental leave can only exist in couples where the mother is willing to give up some leave, and the father willing to take some. Furthermore, Figure 1 shows that for couples where the mother sets a T_{min} to ten weeks, the fathers' preferences is irrelevant, and for fathers' with a threat point T_{max} of zero, the preference of the mother is irrelevant.

Figure 2: Gendered spheres bargaining, with bargaining costs

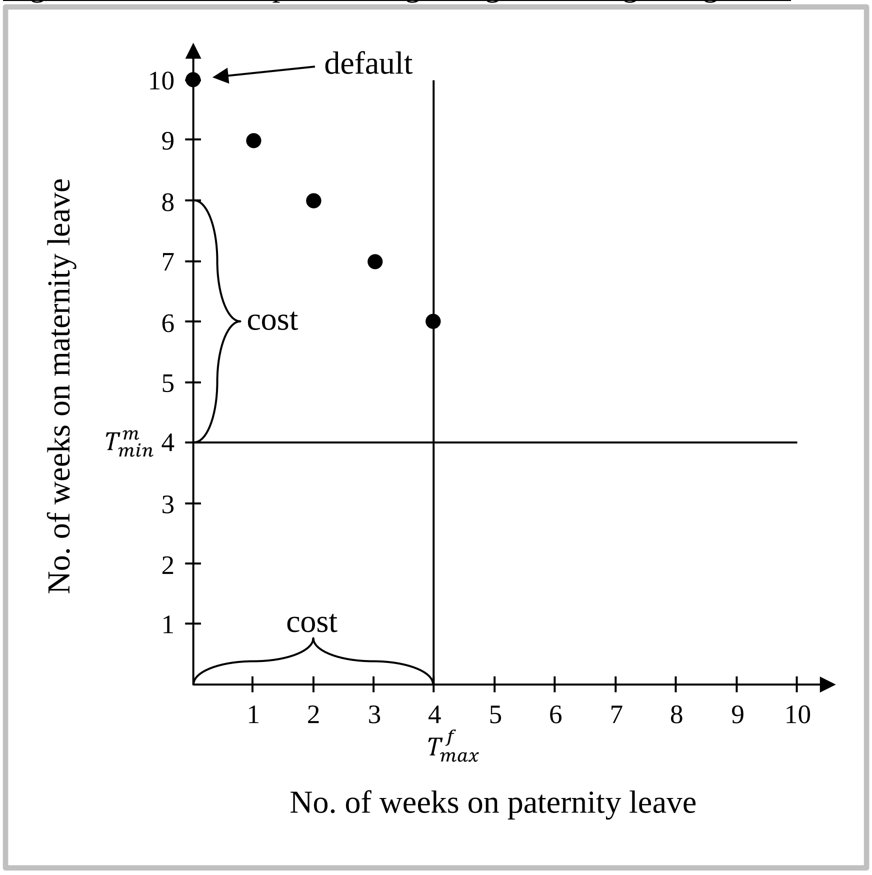

For some couples with high bargaining costs, the gendered spheres push boundaries for cooperation even further (Figure 2). Because of the high cost of engaging in conversations over, planning of and execution of potential leave for them, they may decide against cooperation. Instead of cooperating, the spouses resort to their gendered spheres to a larger extent than they would have without these transaction costs. In the case of Figure 1, the bargaining costs push the threat points over the limit for the father, so he sets his threat point to zero weeks. The spouses only discuss the division of leave if the expected benefits from bargaining exceed the cost of bargaining. If they expect high costs or low gains from discussion altogether, they avoid the discussion, and it becomes optimal for them to remain gender specialized. For example, if the father prefers a maximum of 4 weeks of parental leave, he might decide against bargaining over the weeks, leaving the entire leave to the mother, because the bargaining costs exceed the potential gains.

2.4 The father-quota system with gendered spheres

To minimize the influence of gendered spheres on the take-up and avoid the impact of maternity leave on women's careers, Norway, Sweden, Finland and Iceland have introduced a quota regime that earmarks a minimum amount of weeks for the father. Should he decide not to take up all the earmarked weeks, the mother cannot use them. The couple can, however, decide that the father takes up more of the rest of the leave that the couple shares. Figure 3 shows the potential outcomes of quotas of five earmarked weeks for the father. Due to gendered spheres, the same threat points apply as in the free-choice system, T_{min} for the mother, and T_{max} for the father, but instead of a gendered default of all the parental leave for the mother and none for the father, the new quota default is five weeks for the father and five for the mother.

Figure 3: Quota regime with gendered spheres

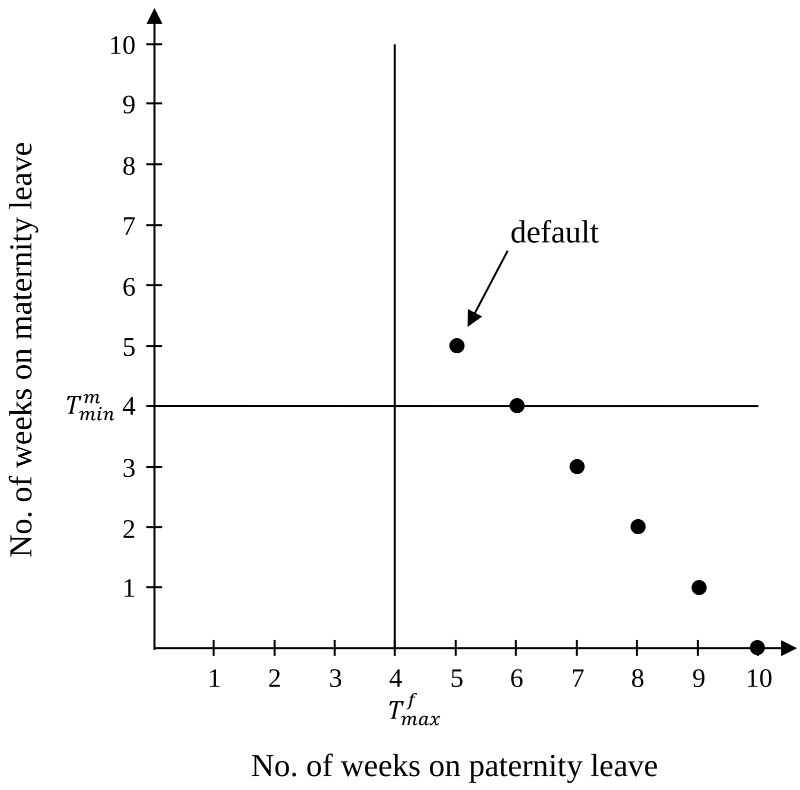

Firstly, the quota regime obviously does not allow for mothers taking up all the parental leave, even when both spouses agree on this division. This undermines the efficiency in families who agrees on a 10/0 split. Secondly, as the gendered spheres are still present, the mother sets a minimum of leave she is willing to give up, and the father a maximum of leave he is willing to take, just as in the free-choice regime. The difference between Figure 1 and Figure 3 is that the possible outcomes have shifted from the mother to the father. However, the threat points remain gendered, and introducing a bargaining cost has the same impact as in Figure 1, namely pushing the leave towards the mother. For most couples in the gendered spheres bargaining set-up, this will be suboptimal, and in most cases fathers will take up the maximum of five weeks, but no more.

Despite the larger take-up of fathers' parental leave in a quota-system, Ekberg, Eriksson and Friebel (2013) find no increase in paternal household work after the parental leave. The presence of gendered spheres could be a potential explanation for the absence of an effect on household work. The spheres still pushes the household work towards the mother, and even if fathers take-up the earmarked part of leave, this does not necessarily change other gendered chores, such as temporary leave when children are sick.

Neither a free-choice nor a quota system makes the gendered spheres disappear. In a dynamic setting with learning over time (for instance following Baysian learning as in Breen and García-Peñalosa (2002)), increasing the number of fathers on leave, even if inefficient for some families, could diminish the gendered spheres over time. For some couples, fathers taking up more leave due to the quota-system may realize that they like being on leave, or mothers that the father can do as good a job as the mother, and in these cases the new insight might influence their take-up of leave for child number two. The notion that parental leave belongs to the mother may diminish. Moreover, if one more week on leave means different marginal costs for each parent – less for the mother because she has already taken up 14 weeks – the number of weeks on earmarked father leave may induce higher costs on fathers per se. However, earmarking would decrease the marginal cost of one extra week. After a period of five weeks on leave, the marginal cost of an extra week may already be decreasing and an extra week on leave will be less costly for the father than if he had only taken up two weeks.

3. Empirical strategy

In this section, I empirically test whether high bargaining costs decrease fathers' take-up of leave, as predicted by theory in a world with gendered spheres. First, I review the data sets in view of the Danish parental leave institutional setting. Second, I define two variables that proxy the cost of bargaining. Third and last, I describe the several covariates that I include in the estimations.

3.1 Data and the Danish institutional setting

For the empirical testing, I combine the Danish longitudinal cohort study of children born 1995 (DALSC) with high-quality register data maintained by Statistics Denmark. Even though the Danish parental leave system has changed considerably since 1996, for international comparisons outside Scandinavia, the 1996 system serves as an attainable case for other countries to follow. [1]

In 1984 it became possible for Danish parents to share parental leave. Two weeks were earmarked for the father (paternity leave), and if the father did not take up the leave within the first 14 weeks of the child's life, the family would lose those weeks. Moreover, the first 14 weeks were

[1] In the 2015 Danish leave system, parents can divide 32 of a total 52 weeks between them. In this paper, I examine an earlier leave system in which the parents could share 10 weeks of a total of 24 weeks.

earmarked for the mother. After these 14 weeks, the couple could divide 10 weeks. When comparing systems across countries, one must be cautious, because adoption can take a long time to increase to high levels (Dahl, Løken, and Mogstad, 2014). In case of father's parental leave, however, the introduction of quotas for fathers in Norway increased the take-up rate for fathers from 3 percent to 35 percent immediately. In 2006 the take-up rate climbed to 70 percent levels, partially due to peer-effects (Dahl, Løken, and Mogstad, 2014).

In the year of the survey in 1996, twelve years after the introduction of the rights to parental leave for fathers, mothers were entitled to full salary during the first 14 weeks. In the public sector the parents received a full salary during leave, whereas the private sector only offered full salary to the mother during the first 14 weeks. Hereafter the parent on leave received wage-replacement of up to 90 percent of the salary, with an upper limit. In 1988 the limit was 88 percent of the average income for mothers, and 55 percent the average income for fathers as fathers averagely earns more (Christoffersen, 1990).

Furthermore, the mother is required to inform her employer about her pregnancy within 13 weeks before her due date. She must also inform her employer whether she plans to take-up the first 14 weeks of leave. The father is required to inform his employer at a later date, within the last four weeks before her due date. They are both required to inform their employers about the length of their respective parental leave before the child is eight weeks old. The information for the employer differs therefore considerable. Fathers who do not wish to take up leave, neither two weeks of father leave nor parental leave, do not have to inform their employers about their new responsibilities, whereas mothers earlier than the father, and in all cases, must inform her employer.

In reality both parents can be on leave at the same time, and DALSC does not distinguish between the two scenarios. Theoretically the gendered spheres assumes separate leave spells, in which case the mother *or* the father takes up the chore of child rearing. If they both take up leave, the mother may still tend to her gendered sphere of child rearing while the father is still tending to his non-child rearing sphere. When setting up the empirical model, I interpret leave as used separately.

I use a Linear Probability Model (LPM) for all analyses. Each family must decide whether the father takes up parental leave[2]:

$$h_m = \begin{cases} 1 \text{ if the father takes up leave} \\ 0 \text{ otherwise} \end{cases}$$

In the analyses, the fathers' uptake of parental leave h_m is the outcome, and the variables of interest are different measures of bargaining costs in the family. To account for different threat points in different segments of society, I control for relative education, relative disposable income, and relative age of the couple. As the absolute measures are also important for the threat point, I include absolute values for the mothers' education, income, and age. Because of collinearity reasons, I cannot include both the mothers' and the fathers' absolute values along with their relative values, which is why I only include variables for the mothers' socioeconomic background.

Controlling for different covariates, I account for some variation that might simultaneously influence the costs of bargaining and fathers' parental leave. These associations between variables remain descriptive, however, as I cannot entirely control for omitted variables and address the issue of reverse causality. The timing of the survey and parental leave is such that parents are interviewed in the middle of or after the potentially shared leave. Consequently I cannot empirically dismiss the possibility that fathers leave directly affect the bargaining costs in the family. Thus the results are purely indicators of how the costs of bargaining and parental leave sharing interact, and how the empirical findings support the theory of bargaining costs.

In DALSC, each parent reports how much parental leave he or she individually used or planned to use. Table 1 shows the distribution of parental leave for fathers and mothers. It also shows that most fathers take up the two weeks of paternity leave, which the couple does not compete for (77 percent took up the two weeks), but when the couples have to decide on parental leave, mothers take up the most (only 5.8 percent of fathers take up parental leave).

[2] For the purpose of the empirical analyses, I can only examine the extensive margins: whether fathers take up leave. How much leave the father takes up is equally interesting, but the data does not provide enough variation for meaningful analyses on the intensive margin.

Table 1: The take-up of earmarked paternity leave, and parental leave for fathers and mothers

Earmarked paternity leave		Parental leave		Parental leave	
Weeks	Fathers	Weeks	Fathers	Weeks	Mothers
0	730	0	3,749	0	21
1	179	1-2	41	1-2	8
2	3,076	3-4	26	3-4	7
		5-6	23	5-10	7
		7-8	17	11-12	8
		9-10	126	13-14	40
				15-16	33
				17-18	23
				19-20	65
				21-22	41
				23-24	3,805

Notes: The intervals differ to make sure that each cell contains more than five individuals to obtain anonymity, as demanded by Statistics Denmark.

3.2 Costs of bargaining

The covariate of interest is the cost of bargaining. Instead of using one measure only in the estimations, I examine two, each of which capture different aspects of costs. The empirical results will shed light on which measure represents the cost of bargaining better. The two measures are:

1) *How often do you fight about daily chores?* Dummy equals 1 if one or both parents answer: "Daily" or "A few times a week". Dummy equals zero if both parents answer one of the other two: "A few times a month" or "Rarely/never".

2) *To what extent do you and your spouse/partner agree on sharing the daily chores?* Dummy equals 1 if one or both parents answer: "Do not really agree" or "Do not agree at all". Dummy equals zero if both parents answer one of the following: "Completely agree" or "Somewhat agree".

In 1) at least one spouse answers that the couple often fights. Women who desire less leave or men who desire more may become discouraged from negotiating altogether, because doing so will result in yet another fight. In 2) couples generally disagree about the daily chores. The qualitative difference from 1) is that the couples do not necessarily fight because of the disagreement, but they know that they have different preferences. Since they know that they most likely disagree if they do discuss different issues, they may choose to avoid discussions.

As a counterfactual to experiencing high bargaining costs, I define a third variable below.

3) *How well do you tackle the situation following a fight?* Dummy equals 1 if one or both parents answer: “We never fight”. Dummy equals zero otherwise, that is if both parents answer one of the following: “Unfortunately we handle fights very badly”, “We become friends again right away”, “We tackle it with humor”, or “After a while it is all forgotten”.

In 3) couples state that they never fight. Never fighting is ambiguous as to costs, as never fighting could be a sign not only of complete cooperation, which would result in a perfect allegiance of preference and bargaining, but never fighting could also mean complete avoidance of conflict, which would result in amplified gendered spheres. This serves as a lower bound estimate for conflict, because never fighting is in most cases better than fighting often or disagreeing daily.

3.3 Controls

I set up a model that includes several measures for the threat point. Like Amilon (2007), I include the mother’s age, education level, and disposable income in 1994 (the year before the baby was born). Furthermore, I include the number of children (if any) prior to the child being born in 1995. I also include a variable for the number of joint children within the current couple. I use the measure for joint children and children from previous relationships as a proxy for how experienced the parents are. In Amilon (2007), the covariates control for the family’s preferences for temporary leave and paid work. As I use a gendered spheres threat point, the threat point in Amilon (2007) is different from mine. Even though the outcome of potential non-cooperation differs, the variables are the same for reaching that threat point.

As Amilon (2007), I include relative measures to capture mothers’ bargaining power. Furthermore, I use the relative measures to grasp the gendered threat point. The gendered spheres are also part of a relative measure between the parents. Therefore, as Amilon (2007), I include several relative measures: education, disposable income and age. Bertocchi and Torricelli (2014) find that these three measures, among others, increases the probability of the woman being in charge in the relationship, in this case, to be responsible for the financial and economic choices in the family.

Datta Gupta and Stratton (2010) show that relative education is a strong predictor of leisure activities, and as such serves as a good measure of bargaining power. In the estimation, relative education is split in three: Either the mother has a higher education level (more than one level difference), they have equal education levels (1 or no difference), or the father has a higher education level (more than one level difference).

Table 2 reports the mean statistics for mothers and fathers in the survey. The table shows fathers having longer educations (9.4 percent vs. 6.4 percent with more than 3 years at the university). The percentage of couples in which the woman has the longest education equals those couples in which the man has the longest education: 18 percent in each group. Men receive higher disposable incomes (17,794 DKK difference), are 2.3 years older, and have a few more children from previous relationships (0.147 children vs. 0.133 children). On average, the couples have 0.63 children with each other before the child is born in 1995.

Table 2: Mean statistics and standard errors for mothers and father, and jointly.

Variables	Mean	SE
Share of fathers on parental leave	0.06	0.23
Mothers' characteristics		
Education level		
Primary school	0.21	0.41
High school	0.13	0.34
Vocational education	0.37	0.48
Bachelor	0.04	0.19
Short and medium term	0.17	0.38
Master degree or higher	0.07	0.25
Missing education information	0.01	0.10
Age	29	4.47
Number of children from previous relationships	0.13	0.47
Disposable income	104,262	33,189
Fathers' characteristics		
Education level		
Primary school	0.21	0.40
High school	0.08	0.27
Vocational education	0.43	0.49
Bachelor	0.07	0.25
Short and medium term	0.10	0.31
Master degree or higher	0.09	0.29
Missing education information	0.02	0.15
Age	32	5.31
Number of children from previous relationships	0.15	0.50
Disposable income	122,138	104,593
Within couples differentials		
Relative education		
Mother has higher education than father (+2 or more levels)	0.19	0.39
Mother has same education as father (+- 1 level)	0.60	0.49
Father has higher education than mother (+2 or more levels)	0.18	0.39
Difference in disposable income (fathers' income - mothers' income)	-17,795	107,821
Age difference	2.35	4.02
Joint children within the couple	0.63	0.79
Bargaining cost dummies		
Cost of bargaining measures		
1) At least one spouse answers that the couple *fight over* daily chores daily or weekly	0.24	0.43
2) At least one spouse answers that the couple *disagree* about daily chores	0.06	0.24
3) At least one spouse answers that the couple *never fights*	0.18	0.38
Observations	3,985	

4. The presence of gendered spheres and bargaining costs

Using the dummy for fathers' parental leave as the outcome, combined with the different bargaining cost measures and the covariates, I estimate the probability that fathers take up parental leave. The findings in Table 3 indicate that the costs of bargaining do influence the division of leave between the parents.

In column 1, including the first measure for bargaining costs – arguing about daily chores – the sign of the estimate is negative and bordering on significant at an 11 percent level. In column 2, comprising the second measure – disagreeing about daily chores – *significantly* decreases the fathers' likelihood of taking up leave. In households with these types of high bargaining costs, there is a decrease in the probability that fathers take up parental leave of 3.2 percentage points. As only 5.8 percent of the fathers in the sample are on parental leave, the decrease from bargaining costs cuts the likelihood of taking up leave in more than half.

In column 3, the third measure – never fighting – neither decreases nor increases fathers' take-up of parental leave significantly. Never fighting seems to be a good indicator of a neutral measure for costs of bargaining, as the estimate is close to zero and insignificant. This should not be interpreted as the non-existence of gendered spheres, but as a measure of low bargaining costs.

The estimates in Table 3 are given in a free-choice system. As the gendered spheres are still present in the quota system, I speculate that the results would be similar, even though the correlation would be smaller because of the higher default of five weeks, as opposed to zero in the free-choice system.

Table 3: The impact of costs of bargaining on the likelihood of fathers' parental leave.

	1		2		3	
	Coef.	SE	Coef.	SE	Coef.	SE
Constant	-0.045	0.031	-0.046	0.031	-0.046	0.030
Education						
Mother's primary school	-0.007	0.012	-0.007	0.012	-0.008	0.012
Mother's high school	0.018	0.012	0.017	0.012	0.012	0.012
Mother's vocational education (reference)						
Mother's bachelor	0.007	0.020	0.006	0.020	0.006	0.020
Mother's short and medium term	0.032***	0.011	0.031***	0.011	0.032***	0.011
Mother's master or higher	0.053***	0.016	0.052***	0.015	0.053***	0.015
Mother's education missing	0.061	0.038	0.059	0.038	0.061	0.037
Mother has higher education than father (+2 or more levels)	-0.021**	0.010	-0.021**	0.010	-0.021**	0.010
Mother has same education as father (+- 1 level) (Reference)						
Father has higher education than mother (+2 or more levels)	0.016	0.011	0.015	0.011	0.016	0.011
Income						
Mother's disposable income	-0.005	0.133	-0.005	0.133	-0.005	0.001
Difference in disposable income (father's income – mother's income)	0.019	0.035	0,020	0,035	0.020	0.034
Age						
Age difference	0.003**	0.001	0.003**	0.001	0.003**	0.001
Mother's age	0.003***	0.001	0.003***	0.001	0.003***	0.001
Children						
Father's no. of children from previous relationships	0.005	0.008	0.005	0.008	0.005	0.008
Mother's no. of children from previous relationships	-0.011	0.009	-0.011	0.009	-0.011	0.009
Joint children	-0.007	0.005	-0.007	0.005	-0.007	0.005
Costs of bargaining						
Couple fight about chores	-0.014	0.008				
Couple disagrees about daily chores			-0.032**	0.015		
Couple never fights					0.0001	0.009
No. of observations	3,985		3,985		3,985	
R^2	0.0134		0.0140		0.0132	

Notes: Standard errors next to the coefficient estimate. *significant at 10%, **significant at 5%, ***significant at 1%

4.2 Relative bargaining power

The gendered spheres are always present in both the free-choice system and the quota-system, but the impact of bargaining costs may be more pronounced in the free-choice system. For spouses with high bargaining power, the gains from initiating an argument over parental leave will, all else equal, be higher, because they are more likely to win the argument. Datta Gupta and Stratton (2010) compare bargaining power measures for U.S. men and women to those of Danish and find that relative education is the best predictor of time spent on leisure activities in both the U.S. and Denmark. Relative education serve as a much better proxy for intra-household bargainigng power than for instance relative income, as income not only proxy bargaining power, but in dynamic models also the outcome of bargaining power. It is less likely that couples increase education as a direct consequence of intra-household bargaining. Consequently I use relative education as a measure for bargaining power. I estimate the same model as in 4.1, the only difference being that I split the samples in three types of couples: mother has the highest education, mother and father have equally long educations, and father has the highest education.

Table 4 shows point estimates for the different cost measures for subsamples of different bargaining power. As I expected, couples in which the mother has the highest education, the estimates are close to zero for all cost measures. The parental leave belongs to her, if she wants, but if she prefers a different distribution, high bargaining costs are overcome by high bargaining power. Table 4 shows similar results for fathers with high bargaining power. He abstains from parental leave by default, but if he prefers some leave after all, the costs do not deter him from bargaining more than the gendered spheres already do.

The largest and most significant correlations between fathers' leave and costs occur for couples with equal bargaining power. The more egalitarian the couple, the larger impact of the costs. This is because the gains from bargaining become smaller, as each spouse is less likely to win the argument. If couples disagree about daily chores, paternal take-up decreases by 4.2 percentage points, which is a large decrease from the mean of 5.8 percent. If they argue about the chores, the point estimate is half that size. In couples who never fight, the father is more likely to take up leave.

Table 4: The impact of bargaining costs on the likelihood of fathers' parental leave, by bargaining power.

Costs of bargaining	Mother longest education	Equally long education	Father longest education
Couple argues about chores	-0.008	-0.020*	-0.003
	(0.018)	(0.011)	(0.021)
Couple disagrees about daily chores	0.002	-0.042**	-0.038
	(0.033)	(0.019)	(0.036)
Couple never fights	0.017	0.008	0.020
	(0.021)	(0.012)	(0.022)
No. of observations	1,805	1,274	896

Notes: Standard errors in parentheses *significant at 10%, **significant at 5%, ***significant at 1%. All the covariates in Table 3 are included in the estimations.

Table 4 also shows that for couples in which the father has the highest bargaining power, arguing about daily chores has a small insignificant impact on the father's leave, whereas disagreeing about chores has a larger, but still insignificant, impact (-3.8 percentage points). As for the other types of families, among couples who never fight the father is more likely to take up leave.

The results in Table 4 suggest that the impact of bargaining costs on fathers' parental leave is mainly driven by couples where bargaining power is equal or highest for the father.

4.3 Experienced parents

Learning is an essential part of parenting. As a first-time parent, the child-rearing obligations can be very different from the prior expectations and beliefs about how well a father can rear children – both from the point of view of the mother and the father – and the expectations may be gender biased. In a free-choice system, fathers take up very little leave due to the gendered spheres, and the learning process will be slow. Even in a quota system the adoption rates for paternity leave are long ((Dahl, Løken, and Mogstad, 2014). For parents with previous children, the expectations will naturally be different.

The gendered spheres in a free-choice system can explain why the estimate of bargaining costs is insignificant for first-time parents, but significant for experienced parents in Table 5. First-time parents may fall back into gendered spheres and not bargain over leave, but for the second time parents, their expectations have been updated. Fathers may realize that they enjoy child-rearing more than they expected to before the first time, whereas the opposite may hold true for mothers. This may lead to sharper differences in preference than the first time, leading to bargaining and its associated costs. If the couple experience high costs, the father takes up even less leave than he did the first time.

Table 5: The impact of costs of bargaining on the likelihood of fathers' parental leave by prior children.

Costs of bargaining	No prior children	At least one spouse has prior children
Couple argues about chores	-0.010	-0.015
	(0.011)	(0.013)
Couple disagrees about daily chores	-0.027	-0.044*
	(0.019)	(0.024)
Couple never fights	-0.007	0.011
	(0.013)	(0.014)
No. of observations	2,320	1,665

Notes: Standard errors in parentheses *significant at 10%, **significant at 5%, ***significant at 1%. All the covariates in Table 3 are included in the estimations.

The result of bargaining costs for experienced parents is likely to be different in a quota-system. Parents update their beliefs after the first child, and as the fathers also take up parental leave, their updated beliefs are different in a quota-system than in a free-choice system. Both parents update their beliefs about their own and the other's abilities much faster in a quota-system. This could potentially decrease the influence of gendered spheres for the second child. If the gendered spheres are less dominant, the influence of bargaining costs will also decrease, as the costs will impact the division of parental leave in a less gendered way.

The impact of gendered spheres is likely to diminish faster over time in a quota-system than in a free-choice system.

5. Conclusion

This paper discusses the theoretical bargaining models that are used in the literature and how their assumptions differ. The unitary model by Becker views the family as one joint unit that optimizes as such, without regard to who earns the income. The empirical literature has discarded the model as being overly simple, and two other types of models have emerged: the cooperative and non-cooperative. The cooperative models assume efficiency within marriage, and spouses optimize a joint utility function. These models assume that the couples have no transaction costs from bargaining, and the one with the most bargaining power obtains the majority of the utility. The spouses have threat points that amount to the utility of their outside options. Each spouse must be at least as happy within marriage as outside the marriage. If that is not the case, they divorce, and the game ends. The beauty of the cooperative models is that they produce efficient outcomes, and thus the differences in outcomes between spouses can be interpreted as a measure for bargaining power.

The non-cooperative models argue that sometimes marriage is inefficient. Couples' counseling and domestic violence prove that we have to consider this option. Instead of optimizing jointly, couples optimize separately, but still consume the public goods jointly. The costs of bargaining – transaction costs – such as sulking, retaliation, and loud arguments dissuade some couples from bargaining, which can explain the separate optimization. If the potential costs are larger than the potential gains from bargaining, the couple refrains from doing so, but continues to benefit from marriage through joint consumption. Their threat points can be divorce as in the cooperative models, but they can also be less severe.

In this paper, I argue that another threat point is more likely. When spouses have a child and decide how to divide their parental leave, they are not likely to divorce. In this rare game (which is unlikely to occur more than a few times in life) non-cooperation between spouses is more likely to result in gendered spheres. Lundberg and Pollak's (1993) gendered spheres bargaining model states that couples have different spheres for different types of work. In this case, parental leave is within the mothers' gendered sphere and paid work within the fathers'. I argue that the gendered spheres bargaining model is the best fit for discussing intra-household division of parental leave. Due to gendered spheres the parental leave by default belongs to the mother. The only room for bargaining over leave is for couples where the mother is willing to give up some leave, and the father is willing to take up some. If this is not the case, the space for negotiation disappears. The gendered spheres are a likely reason for why fathers in Denmark, a country in which couples can freely decide on the division, take up only 7.4 percent of the entire parental leave (Nordic Statistical Yearbook, 2012).
Comparing the free-choice system to the father-quota system my conclusion is that gendered spheres are still present in a quota-system, but the impact of gendered spheres is likely to diminish faster in quota-system than in a free-choice system.

Empirically I use The Danish Longitudinal Study of the 1995 cohort to examine the presence of gendered spheres in the Danish free-choice system. Using a linear probability model with several covariates for family type, and thus different gender stereotypes, I estimate the correlation between costs of bargaining and fathers' take-up of parental leave, and find that bargaining costs significantly correlates negatively with paternal parental leave. Performing two subsample analyses, I find that the correlation between bargaining costs and fathers' parental leave is strongest among couples with equal bargaining power and couples with previous children. The empirical findings point to the presence of gendered spheres. Because the fathers' in high costs families are less likely to take up leave, the gendered spheres are likely to be present. Because the available data cannot fully account for omitted variables bias and reverse causality, the empirical findings are purely descriptive and more empirical research is needed to fully understand the interactions between bargaining costs and fathers' take up of parental leave. The empirical findings however correspond with the theoretical framework of bargaining costs.

The findings in this article illustrate and expand on the political debate over leave. Instead of discussing incentives for fathers or mothers only, politicians ought to discuss how to optimize the parental leave system to support cooperation among spouses, rather than supporting the gendered status quo – that women take up the lion's share of parental leave. Earmarked paternal leave has proven successful in increasing fathers' take-up, but we nonetheless need more research in how to counteract the gendered spheres and thus the costs of bargaining to nudge couples towards cooperative, and less gendered outcomes.

Acknowledgements I wish to thank my supervisors at Aarhus University Nabanita Datta Gupta and Marianne Simonsen, my supervisor at SFI, Beatrice Schindler-Rangvid, AU RECEIV center (project no. 908792) and SFI. I gratefully acknowledge the valuable comments from the participants at the RECEIV workshop 2015, from Shelly Lundberg and Anna Amilon.

References

Amilon, A. (2007). On the Sharing of Temporary Parental Leave: The Case of Sweden, *Review of Economics of the Household* 5(4): 385–404

Becker, G. (1981). A Treatise on the Family. Cambridge, Mass.: Harvard University Press. 1st Edition 1981, Enlarged Edition 1991.

Berger, L. M., Hill, J., and Waldfogel, J. (2005). Maternity leave, early maternal employment and child health and development in the US*. *The Economic Journal*, *115*(501), F29-F47.

Bertocchi, G., Brunetti, M. and Torricelli, C. (2014). Who holds the purse strings within the household? The determinants of intra-family decision making. *Journal of Economic Behavior & Organization*, *101*, 65-86.

Boll, C., Leppin, J., and Reich, N. (2014). Paternal childcare and parental leave policies: evidence from industrialized countries. *Review of Economics of the Household*, *12*(1), 129-158.

Breen, R. and García-Peñalosa, C. (2002) Bayesian learning and gender segregation. *Journal of Labor Economics*, *20*(4), 899-922.

Browning, M. and P.-A. Chiappori (1998). Efficient Intra-Household Allocations: A General Characterization and Empirical Tests. *Econometrica* 66(6),1241- 1278.

Browning, M., and Gørtz, M. (2012). Spending Time and Money within the Household. *The Scandinavian Journal of Economics*, *114*(3), 681-704.

Chiappori, P.-A. (1988). Rational Household Labor Supply. Econometrica 56(1), pp. 63-90.

Chiappori, P.-A. (1992). Collective Labor Supply and Welfare. *Journal of Political Economy* 100(3), pp. 437-467.

Chiappori, P.-A. (1997). Introducing Household Production in Collective Models of Labor Supply. *Journal of Political Economy* 105(1), pp. 191-209.

Christoffersen, M. (1990). Maternity leave - men's and women's professional background to take leave.[*Barselsorlov - mænds og kvinders erhvervsmæssige baggrund for at tage orlov]*. Copenhagen: SFI 90:18.

Cools, S., Fiva J. H., and Kirkebøen L. J. (2015). Causal effects of paternity leave on children and parents. *The Scandinavian Journal of Economics*, 117, 801–828.

Datta Gupta, N. and Stratton, L. S. (2010). Examining the impact of alternative power measures on individual time use in American and Danish couple households. *Review of Economics of the Household*, *8*(3), 325-343.

Drange, N. and Rege M. (2013). Trapped at home: The effect of mothers' temporary labor market exit on their subsequent work career. *Labour Economics*, *24*(10), 125-136.

Ekberg, J., Eriksson, R., and Friebel, G. (2013). Parental leave—A policy evaluation of the Swedish "Daddy-Month" reform. *Journal of Public Economics,97*, 131-143.

Ejrnæs, M., and Kunze A. (2013). Work and wage dynamics around childbirth. *The Scandinavian Journal of Economics*, 115(3), 856-877.

Haagensen, K. M. (Ed.). (2012). *Nordic Statistical Yearbook 2012*. Nordic Council of Ministers.

Hardoy, I., and Schøne, P. (2013). Enticing even higher female labor supply: The impact of cheaper day care. *Review of Economics of the Household*, 1-22.

Konrad, K. A., and Lommerud, K. E. (2000). The bargaining family revisited.*Canadian Journal of Economics/Revue canadienne d'économique*, *33*(2), 471-487.

Lundberg, S., and Pollak, R. A. (1993). Separate spheres bargaining and the marriage market. *Journal of political Economy*, 988-1010.

Lundberg, S. J., Pollak, R. A., and Wales, T. J. (1997). Do husbands and wives pool their resources? Evidence from the United Kingdom child benefit. *Journal of human resources*, 463-80.

Lundberg, S., and Pollak, R. A. (2003). Efficiency in marriage. *Review of Economics of the Household*, *1*(3), 153-167.

Nielsen, H. S. (2009). Causes and Consequences of a Father's Child Leave: Evidence from a Reform of Leave Schemes. *IZA discussion paper* no. 4267, Institute for the Study of Labour, Bonn, Gernmany

Rasmussen, A. W. (2010). Increasing the length of parents' birth-related leave: The effect on children's long-term educational outcomes. *Labour Economics*,*17*(1), 91-100.

Simonsen, M. (2010). Price of High-quality Daycare and Female Employment.*The Scandinavian Journal of Economics*, *112*(3), 570-594.

Waldfogel, J. (1998). Understanding the" family gap" in pay for women with children. *The Journal of Economic Perspectives*, 137-156.

DEPARTMENT OF ECONOMICS AND BUSINESS ECONOMICS
AARHUS UNIVERSITY
SCHOOL OF BUSINESS AND SOCIAL SCIENCES
www.econ.au.dk

PhD dissertations since 1 July 2011

2011-4	Anders Bredahl Kock: Forecasting and Oracle Efficient Econometrics
2011-5	Christian Bach: The Game of Risk
2011-6	Stefan Holst Bache: Quantile Regression: Three Econometric Studies
2011:12	Bisheng Du: Essays on Advance Demand Information, Prioritization and Real Options in Inventory Management
2011:13	Christian Gormsen Schmidt: Exploring the Barriers to Globalization
2011:16	Dewi Fitriasari: Analyses of Social and Environmental Reporting as a Practice of Accountability to Stakeholders
2011:22	Sanne Hiller: Essays on International Trade and Migration: Firm Behavior, Networks and Barriers to Trade
2012-1	Johannes Tang Kristensen: From Determinants of Low Birthweight to Factor-Based Macroeconomic Forecasting
2012-2	Karina Hjortshøj Kjeldsen: Routing and Scheduling in Liner Shipping
2012-3	Soheil Abginehchi: Essays on Inventory Control in Presence of Multiple Sourcing
2012-4	Zhenjiang Qin: Essays on Heterogeneous Beliefs, Public Information, and Asset Pricing
2012-5	Lasse Frisgaard Gunnersen: Income Redistribution Policies
2012-6	Miriam Wüst: Essays on early investments in child health
2012-7	Yukai Yang: Modelling Nonlinear Vector Economic Time Series
2012-8	Lene Kjærsgaard: Empirical Essays of Active Labor Market Policy on Employment
2012-9	Henrik Nørholm: Structured Retail Products and Return Predictability
2012-10	Signe Frederiksen: Empirical Essays on Placements in Outside Home Care
2012-11	Mateusz P. Dziubinski: Essays on Financial Econometrics and Derivatives Pricing

2012-12 Jens Riis Andersen: Option Games under Incomplete Information

2012-13 Margit Malmmose: The Role of Management Accounting in New Public Management Reforms: Implications in a Socio-Political Health Care Context

2012-14 Laurent Callot: Large Panels and High-dimensional VAR

2012-15 Christian Rix-Nielsen: Strategic Investment

2013-1 Kenneth Lykke Sørensen: Essays on Wage Determination

2013-2 Tue Rauff Lind Christensen: Network Design Problems with Piecewise Linear Cost Functions

2013-3 Dominyka Sakalauskaite: A Challenge for Experts: Auditors, Forensic Specialists and the Detection of Fraud

2013-4 Rune Bysted: Essays on Innovative Work Behavior

2013-5 Mikkel Nørlem Hermansen: Longer Human Lifespan and the Retirement Decision

2013-6 Jannie H.G. Kristoffersen: Empirical Essays on Economics of Education

2013-7 Mark Strøm Kristoffersen: Essays on Economic Policies over the Business Cycle

2013-8 Philipp Meinen: Essays on Firms in International Trade

2013-9 Cédric Gorinas: Essays on Marginalization and Integration of Immigrants and Young Criminals – A Labour Economics Perspective

2013-10 Ina Charlotte Jäkel: Product Quality, Trade Policy, and Voter Preferences: Essays on International Trade

2013-11 Anna Gerstrøm: World Disruption - How Bankers Reconstruct the Financial Crisis: Essays on Interpretation

2013-12 Paola Andrea Barrientos Quiroga: Essays on Development Economics

2013-13 Peter Bodnar: Essays on Warehouse Operations

2013-14 Rune Vammen Lesner: Essays on Determinants of Inequality

2013-15 Peter Arendorf Bache: Firms and International Trade

2013-16 Anders Laugesen: On Complementarities, Heterogeneous Firms, and International Trade

2013-17 Anders Bruun Jonassen: Regression Discontinuity Analyses of the Disincentive Effects of Increasing Social Assistance

2014-1 David Sloth Pedersen: A Journey into the Dark Arts of Quantitative Finance

2014-2 Martin Schultz-Nielsen: Optimal Corporate Investments and Capital Structure

2014-3 Lukas Bach: Routing and Scheduling Problems - Optimization using Exact and Heuristic Methods

2014-4 Tanja Groth: Regulatory impacts in relation to a renewable fuel CHP technology: A financial and socioeconomic analysis

2014-5 Niels Strange Hansen: Forecasting Based on Unobserved Variables

2014-6 Ritwik Banerjee: Economics of Misbehavior

2014-7 Christina Annette Gravert: Giving and Taking – Essays in Experimental Economics

2014-8 Astrid Hanghøj: Papers in purchasing and supply management: A capability-based perspective

2014-9 Nima Nonejad: Essays in Applied Bayesian Particle and Markov Chain Monte Carlo Techniques in Time Series Econometrics

2014-10 Tine L. Mundbjerg Eriksen: Essays on Bullying: an Economist's Perspective

2014-11 Sashka Dimova: Essays on Job Search Assistance

2014-12 Rasmus Tangsgaard Varneskov: Econometric Analysis of Volatility in Financial Additive Noise Models

2015-1 Anne Floor Brix: Estimation of Continuous Time Models Driven by Lévy Processes

2015-2 Kasper Vinther Olesen: Realizing Conditional Distributions and Coherence Across Financial Asset Classes

2015-3 Manuel Sebastian Lukas: Estimation and Model Specification for Econometric Forecasting

2015-4 Sofie Theilade Nyland Brodersen: Essays on Job Search Assistance and Labor Market Outcomes

2015-5 Jesper Nydam Wulff: Empirical Research in Foreign Market Entry Mode

2015-6	Sanni Nørgaard Breining: The Sibling Relationship Dynamics and Spillovers
2015-7	Marie Herly: Empirical Studies of Earnings Quality
2015-8	Stine Ludvig Bech: The Relationship between Caseworkers and Unemployed Workers
2015-9	Kaleb Girma Abreha: Empirical Essays on Heterogeneous Firms and International Trade
2015-10	Jeanne Andersen: Modelling and Optimisation of Renewable Energy Systems
2015-11	Rasmus Landersø: Essays in the Economics of Crime
2015-12	Juan Carlos Parra-Alvarez: Solution Methods and Inference in Continuous-Time Dynamic Equilibrium Economies (with Applications in Asset Pricing and Income Fluctuation Models)
2015-13	Sakshi Girdhar: The Internationalization of Big Accounting Firms and the Implications on their Practices and Structures: An Institutional Analysis
2015-14	Wenjing Wang: Corporate Innovation, R&D Personnel and External Knowledge Utilization
2015-15	Lene Gilje Justesen: Empirical Banking
2015-16	Jonas Maibom: Structural and Empirical Analysis of the Labour Market
2015-17	Sylvanus Kwaku Afesorgbor: Essays on International Economics and Development
2015-18	Orimar Sauri: Lévy Semistationary Models with Applications in Energy Markets
2015-19	Kristine Vasiljeva: Essays on Immigration in a Generous Welfare State
2015-20	Jonas Nygaard Eriksen: Business Cycles and Expected Returns
2015-21	Simon Juul Hviid: Dynamic Models of the Housing Market
2016-1	Silvia Migali: Essays on International Migration: Institutions, Skill Recognition, and the Welfare State
2016-2	Lorenzo Boldrini: Essays on Forecasting with Linear State-Space Systems
2016-3	Palle Sørensen: Financial Frictions, Price Rigidities, and the Business Cycle
2016-4	Camilla Pisani: Volatility and Correlation in Financial Markets: Theoretical Developments and Numerical Analysis

2016-5 Anders Kronborg: Methods and Applications to DSGE Models

2016-6 Morten Visby Krægpøth: Empirical Studies in Economics of Education

2016-7 Anne Odile Peschel: Essays on Implicit Information Processing at the Point of Sale: Evidence from Experiments and Scanner Data Analysis

2016-8 Girum Dagnachew Abate: Essays in Spatial Econometrics

2016-9 Kai Rehwald: Essays in Public Policy Evaluation

2016-10 Reza Pourmoayed: Optimization Methods in a Stochastic Production Environment

2016-11 Sune Lauth Gadegaard: Discrete Location Problems – Theory, Algorithms, and Extensions to Multiple Objectives

2016-12 Lisbeth Palmhøj Nielsen: Empirical Essays on Child Achievement, Maternal Employment, Parental Leave, and Geographic Mobility

ISBN: 9788793195400