NHH

Master Thesis, Economics and Energy, Natural Resources and the Environment

The Effects of Environmental Regulation and Technological Advancement on Labor Demand

Evidence from the Norwegian Primary Aluminum Industry

Tora Lindheim and Nora Austad Sværen
Supervisor: Aline Bütikofer

Norwegian School of Economics
Bergen, Spring 2019

Acknowledgements

We would like to start by thanking our supervisor Aline Bütikofer for advice on the choice of a feasible and interesting topic, and for valuable and constructive inputs on our work throughout the writing process. Her guidance related to econometric analysis has been of inestimable value.

Further, we would like to thank Christian Braathen for helpful insight into LaTeX, and for informative suggestions along the way. We wish to direct thanks to Ingrid Gaarder Harsheim, Anders Halland, Vilde Nakkim, and Shelby Dennis, for their contribution with extensive feedback and proofreading. Lastly, we also wish to thank friends and family for support and care along the way in the time writing our master thesis.

Abstract

The motivation to implement environmental regulations is to increase environmental quality, as well as to promote higher quality of life and health benefits, for citizens. But how do environmental regulations affect labor demand in the regulated industries? The aim of this thesis is to analyze the effect of two environmental regulations imposed on the Norwegian primary aluminum industry on labor demand. Moreover, we investigate the effect of a technological change originating from the latter regulation on labor demand, being the shutdown of the more polluting and less efficient production technology, the Søderberg anodes. In particular, we exploit the introduction of an environmental agreement in 1997, and an increased stringency in the industry's emission permits in 2000, as exogenous sources of variation in labor demand in the pertinent municipalities, by using a differences-in-differences strategy. Furthermore, we instrument the latter regulation on the shutdown of the Søderberg anodes. Using Norwegian municipality–level data, we find that the 97–agreement, and the shutdown of the Søderberg anodes, are associated with a short–term increase in the employment rate equivalent to respectively 135 and 406 jobs, in each of the municipalities in Norway that encompass a primary aluminum production plant.

Acronyms

2SLS Two-Stage Least Squares

CAA Clean Air Act

CAAA Clean Air Act Amendments

CICERO Centre for International Climate and Environmental Research

DiD Differences-in-differences

EPA Environmental Protection Agency

ETS Emission Trading System

EU European Union

IV Instrumental Variables

GHG Greenhouse Gas

LFS Labor Force Survey

MoCE Ministry of Climate and Environment

NEA Norwegian Environment Agency

NSD Norwegian Center for Research Data

NAV Norwegian Labor and Welfare Administration

NOK Norwegian kroner

OLS Ordinary Least Squares

OVB Omitted Variables Bias

SSB Statistics Norway

Contents

List of Figures

List of Tables

1 Introduction

1.1 Motivation and Purpose

The Norwegian Government passed the Pollution Control Act in 1981. It took effect in 1983 to secure environmental quality by reducing existing pollution and the quantity of waste, promoting health benefits, and ensuring sustainable exploitation of nature (MoCE, 1981). Since then, the principal political measures applied to meet the standards of the Pollution Control Act have been taxes and voluntary agreements, in addition to overall measures introduced by the European Union (EU) (Vevatne, Lindhjem, Eskeland, Haugland, & Gullberg, 2004). The manufacturing industry was, until 2007, the largest contributor to environmental pollution in Norway (Miljøstatus, 2018b). Yet thanks to strong pro–manufacturing lobby efforts in the late 1990s, the primary aluminum industry was exempted from the national political measures put in place before 2013 (NEA & SSB, 2013; Kasa, 1999). A voluntary agreement was signed in 1997 by the primary aluminum industry and the Ministry of Climate and Environment (MoCE, Norwegian: "Miljødirektoratet") with a commitment to reduce the emission of greenhouse gases by 2005 (MoCE, 2000, p. 56).[1] Moreover, in 2000, an increased stringency in the emission permits for the primary aluminum industry was introduced (NEA, 2000).

There is a common understanding that environmental regulations lead to a decrease the level of employment at workplaces (Ntb, 2014; Cover, 2011). However, the effect of environmental regulation on labor demand is inconclusive according to theoretical predictions (Berman & Bui, 2001). Benefits of the related reduction of greenhouse gas (GHG) emissions are often defined as the value of the higher environmental quality, potentially leading to improved health and even saved lives. Thus, GHG reduction would be beneficial for the society in the form of the reduced need for health care (Arrow et al., 1996) and a higher quality of life for citizens. However, these benefits entail costs deriving from increased pollution abatement activities for the regulated entities. Such activities will typically demand the installation of pollution abatement equipment, which may be more or less labor intensive than prior to regulation (i.e., the substitution effect) (Berman & Bui, 2001). Alternatively, these activities can lead to increased production costs, causing the plant to reduce its output and reduce demand for labor (i.e., the output effect). It is not given which of the effects will dominate, calling for evidence from empirical studies. If the reduced demand for labor dominates, either frictional, structural, or cyclical unemployment is induced. There are social costs from unemployment, related to the

[1] The MoCE has a particular responsibility for carrying out the environmental policies of the Norwegian government (Government.no, n.d.).

physical output the unemployed could have produced had they not been out of work, as well as other societal costs, including the consequences for the mental and physical health of the unemployed (Moosa, 1997). Investigation of the actual consequences of environmental regulation on labor demand is therefore beneficial for policy–makers when evaluating the implementation of environmental regulation.

The Norwegian primary aluminum industry is a major employer in Norway's districts (Regjeringen, 2000), particularly in seven municipalities located in the southwestern–part of the country. Concerning the agreement made in 1997 (referred to as the 97–agreement) and the subsequent increased stringency of emission permits in 2000 (referred to as the 2000–restriction), the Norwegian media has indicated that this regulation raised unemployment and threatened local communities (Norsk Telegrambyrå, 2003). One particular pollution abatement activity, the shutdown of the Søderberg anodes in the primary aluminum plants located in the municipalities Karmøy, Årdal, Høyanger, Sunndal, and Vefsn, impacted labor demand considerably (Berge, 2008).[2] More precisely, the shutdowns of the Søderberg anodes in Hydro Årdal and Hydro Karmøy were expected to bring about the loss of 500 jobs (Johansen, 2008) and 450 jobs (Norsk Telegrambyrå, 2008), respectively. However, Johansen (2008) indicates that Årdal experienced a successful and positive reallocation of the labor force, and that new businesses were established making use of the municipality's competitive advantages.

Despite the widespread media coverage, there is no empirical evidence exploring the effects of environmental regulation on labor demand in Norway. However, there is extensive international literature on the topic drawing ambiguous conclusions, which is further explored in Section 3. This thesis aims to amplify the reach of the literature by investigating the effects of the increased regulatory stringency related to environmental concerns in the primary aluminum industry, starting with the 97–agreement, and labor demand. In particular, it is interesting to examine the effects in the communities associated to the seven primary aluminum plants. Utilizing panel data, we examine the impact of environmental regulation on employment and unemployment using differences–in–differences (DiD) and instrumental variables (IV) estimations separately. We exploit the two dimensions of variation, across municipalities and time, to estimate the net effects of environmental regulation on labor demand. The dataset constructed, consists of information on the unemployment rate, employment rate, payroll tax paid to the municipality, demographic characteristics, and a geographic characteristic.[3]

[2]The production of aluminum takes place in an electrolytic cell which consists of a carbon anode (Totten & MacKenzie, 2003).

[3]Whether a municipality has a coastline or not.

When applying the DiD approach, the seven aluminum municipalities (constituting the treatment group) are compared with a subset of 144 municipalities (constituting the comparison group) before and after treatment in 1997, as well as after the 2000–restriction. To account for the long compliance period of the 97–agreement and the 2000–restriction, the period following 1997 is divided in two: the years 1997-1999 and the years 2000-2010. Thus, we operate with two cut–offs, 1996/1997 and 1999/2000, to measure the true causal effect on three outcome variables of interest: **the unemployment rate**, **the employment rate**, and **the payroll tax per person in the working age population**. We check the robustness of the results proceeding from our DiD estimation by verifying the fulfillment of the fundamental assumptions of the DiD approach, omitting one of each of the aluminum municipalities singly from the estimations, and lastly, by running a placebo test pretending the treatment happened in 1995 and 1996.

Furthermore, IV estimation is used to explore the direct effect of the shutdown of the Søderberg anodes, a technological advancement, on the same outcomes of labor demand and subset of municipalities. More specifically, we instrument a dummy variable for the aluminum municipalities, in the years posterior to the 2000–restriction on a dummy variable for the years following the shutdown of the Søderberg anodes in five of the seven aluminum municipalities. The IV results are examined by expanding the last year of observations in the sample from 2010 to 2013, and by verifying the fulfillment of the necessary IV assumptions.

Consistent with prior literature (i.e., Berman & Bui, 2001; Ferris, Shadbegian, & Wolverton, 2014), we find evidence that the 97–agreement and the technological advancement, implemented as a result of the 2000–restriction, increases the employment rate in the short–term, and that the combination of the 2000–restriction and the 97–agreement decreases the unemployment rate, employment rate, and the payroll tax per person, after 2000.

1.2 Research Question

Based on the previous subsection, this thesis aims to investigate the following research question:

How did the environmental regulations, namely the 97–agreement and the 2000–restriction, and the related technological advancement, affect local labor demand in the municipalities which encompass a primary aluminum plant?

The remainder of the thesis proceeds as follows. Section 2 provides background information on environmental regulations in Norway and the primary aluminum industry. Section 3 gives an overview of previous literature on environmental regulations and labor demand, on environmental regulations and technological change, as well as on technological change and labor demand. The data applied for our estimations is presented in Section 4, and a thorough description of our empirical approach is provided in Section 5. The presentation of our results takes place in Section 6, before these are discussed and reviewed in Section 7. Lastly, we present our concluding remarks in Section 8.

2 Background

Research has shown that the primary aluminum industry has been exempted from most environmental regulations in Norway in the period from 1991 until its inclusion in the Norwegian Carbon Credit Procurement Program in 2013.[4,5] However, in the absence of a better alternative, a voluntary agreement was signed between the primary aluminum industry and the Norwegian government in 1997. In the following section, we present the evolution of the political agenda for environmental protection within the time period 1992-2010, before we present more closely the characteristics of the primary aluminum industry in Norway, as well as our case of study. Figures 1 and 2 depict a summary of the evolution of the Norwegian political landscape regarding overall environmental considerations, as well as their more specific applications to the primary aluminum industry.

2.1 Environmental Regulations in Norway

The Norwegian MoCE was established in 1972, and was one of the first ministries in the world with particular responsibility for climate and environmental concerns (Reitan, 1998). Despite the early allocation of responsibility for climate and environmental concerns, the idea that environmental policies were in conflict with considerations for economic efficiency and growth emerged in the early 1970s (Reitan, 1998). Thus, employment and output were given higher priority than environmental considerations–not an unexpected standpoint from one of the world's largest petroleum exporters. Surprisingly, Norway subsequently became a leading activist on the international scene for environmental policies in the second part of the 1980s (Andresen & Butenschøn, 2001).

An emphasis on nationwide legally bound environmental protection commenced in October 1983 once the Pollution Control Act took effect. The Act was introduced to retain a reasonable level of environmental quality, by ensuring that emission and waste did not cause damage to human health or adversely affect welfare, or the outdoor environment and its capacity for self-renewal (MoCE, 1981). According to the Pollution Control Act, all emissions of GHGs are considered pollution. As a ground rule, GHG-pollution has from that time onward, been bound by regulation. Hence, all operations needed (and still need) permission by either concession or law in the Act.

[4] "Procurement of carbon credits is a supplement to national measures to reduce global greenhouse gas emissions."(MoCE, 2018)

[5] The production from the aluminum producing companies have been under surveillance through the issuance of emission permits. The first emission permits were issued in the 1970s.

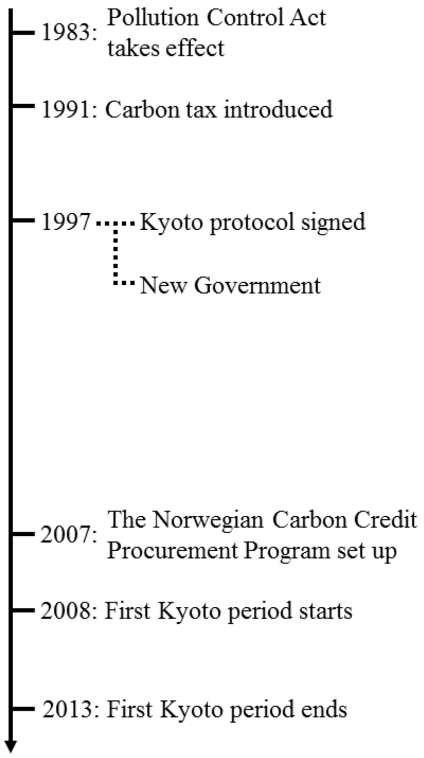

Figure 1: Environmental Regulations in Norway, 1983-2013

Notes: Timeline is based on the information presented in Section 2.1

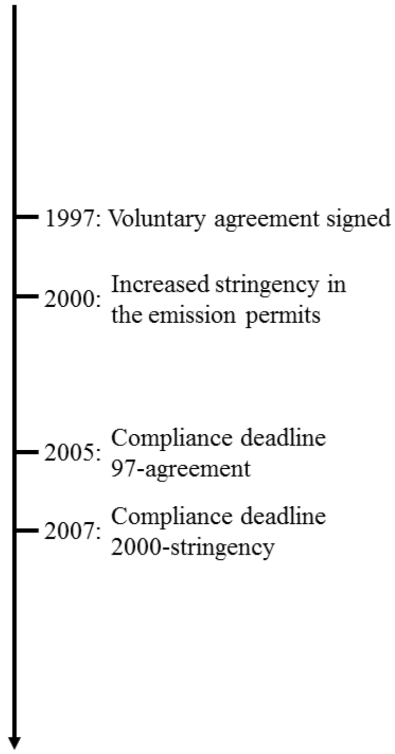

Figure 2: Environmental Regulations in the Norwegian Primary Aluminum Industry, 1983-2013

Notes: Timeline is based on the information presented in Section 2.3

In 1989, the Norwegian parliament agreed to stabilize Norwegian CO_2 emissions at that year's level by 2000 (Reitan, 1998). Not long after, this was questioned by the manufacturing industry and was further quietly abandoned (Andresen & Butenschøn, 2001). As a result, by early 1991, new elements were added to the political agenda. The stabilization of Norwegian CO_2 emissions and taxation was to be both cost–effective and equitable in order to minimize the overall economic losses (Kasa, 1999;2000). Norway was, in 1991, among the first countries to introduce a CO_2–tax (MoCE, 2001). However, energy–intensive industries, like the metallurgical industry, were exempted (Vevatne et al., 2004).[6] According to Kasa (2000), there was intense pro-manufacturing lobbying taking place during this period, which can explain why certain industries were exempted. Later in the 1990s, around the time of the election of a new government, momentum towards increasing and broadening the CO_2–tax to include the exempted industries re–emerged, but it was met with defeat (Kasa, 1999; Alfsen, 1999).

[6]Together with petroleum and transport, the manufacturing industry was one of the three major polluters of GHGs in 2010 (NEA & SSB, 2013).

Until the Climate Convention in Kyoto in 1997, the pro-manufacturing attitude had continued to evolve in Norway. Norway was among the few OECD countries to participate in Kyoto without a domestic target, although it was the first country to adopt one in 1981 (Andresen & Butenschøn, 2001). By the end of the Convention in Kyoto, it was agreed upon that Norway should be allowed emissions one percent above the 1990 level in the first Kyoto commitment period from 2008-2012. The Kyoto protocol was signed 11 December 1997, and it implied an increased focus on the reduction of six GHGs. In accordance with the Kyoto Protocol, Norway was legally bound to take responsibility for its pollution, and to put political measures in place to meet the restrictions of the first commitment period 2008-2012, and the later period 2013-2020 (Alfsen, 1998; MoCE, 2001). In a Government White Paper from the MoCE (2001) published after the Kyoto Convention, Norway's commitment to fighting man-made emissions are emphasized in line with the protocol. It indicates primarily commitment to decrease emissions in 2008-2012, the first Kyoto commitment period, but also commitment to show a demonstrable improvement by 2005 (MoCE, 2001). In 2007, the Norwegian Carbon Credit Procurement Program was set up as an integrated part of the EU Emission Trading System (ETS). The system's purpose was to ensure that Norway would be able to meet its target in the first Kyoto commitment period, yet it excluded the process industry and agriculture industry, which were first included after revision in 2013 (MoCE, 2018; Moen, 2007).[7]

In summary, we have shown in this subsection that while Norway was one of the first to adopt a domestic approach to environmental challenges, the aluminum industry was sheltered from the national measures established between 1991 and 2012, despite being one of the largest polluters. In the rest of this section, we will focus on the primary aluminum industry, and unless otherwise specified, we will refer to the primary production of aluminum as the aluminum industry.

[7]The process industry includes the metallurgical industry (e.g., the aluminum industry), and is a part of the manufacturing industry(Innovation Norway, n.d.). The process industry consists of the "metallurgical industry, pulp and paper, fertilizers, chemicals and pharmaceutical industry stands for half of the export from Norway (not including oil and gas).

2.2 The Primary Aluminum Industry in Norway

The production of primary and secondary aluminum constitutes the largest part of the Norwegian metallurgical industry (Bye, Larsson, & Døhl, 1999). Metal production has long roots in Norway (Godal, 1998). In the aftermath of the Second World War, the Norwegian government took over all enemy properties (Moen, 2007). Among these were the aluminum projects at Årdal and Sunndalsøra, finished respectively in 1948 and 1954. These operations became part of the state-owned company Årdal and Sunndal Verk (Store Norske Leksikon, 2018a). Today, there are seven aluminum plants in Norway, situated in the municipalities Årdal, Høyanger, Kvinnherad, Karmøy, Farsund, Vefsn, and Sunndal. These are operated and owned by Hydro, Alcoa, and Rio Tinto Alcan jointly owns one plant with Hydro (Sør-Norge Aluminum) (The Federation of Norwegian Industries, n.d.). See Table 1 for more information about the start of production and the company operating each aluminum plant.

Table 1: Primary Aluminum Plants in Norway

Plant	Production Start	Company	Share of inhabitants in employable age, 1996	Shutdown Søderberg anodes
Høyanger	1917	Hydro	12.0%	2006
Årdal	1948	Hydro	18.9%	2007
Sunndal	1954	Hydro	18.9%	2002
Karmøy	1967	Hydro	4.5%	2009
Lista (Farsund)	1971	Alcoa	9.2%	New–Søderberg technology
Mosjøen (Vefsn)	1958	Alcoa	7.6%	2003
Husnes (Kvinnherad)	1965	Sør–Norge Aluminum	5.4%	Never had Søderberg technology

Notes: Adopted table from Godal (1998, p. 10). The first column shows plant and location name. Name of the plant and municipality name are equal for the first four, whereas, for the three latter the municipality name is written in parenthesis. The second column shows the year of production start. Information on production start is obtained from Hydro (n.d.-b; n.d.-a; n.d.-d; n.d.-c), Olsen (1996) and Store Norske Leksikon (2018b). The third column shows the names of the plant operators. The fourth column shows the share of inhabitants in employable age in the municipality. The fifth column shows the year of Søderberg–shutdown, and whether a plant did not experience shutdown. Year of shutdown of the Søderberg anodes is obtained from Teknisk Ukeblad (2008), Stavanger Aftenblad (2011) Alcoa (n.d.) and Tjelmeland (1987, p. 115).

Norway is Europe's largest producer of aluminum with a production of about 1.2 million tonnes per year (The Federation of Norwegian Industries, n.d.). The aluminum industry is exposed to international competition (Vevatne et al., 2004, p. 69), and because it is capital intensive (SSB, 2017), large fluctuations in the aluminum price may result in large variations in the companies' net income (Godal, 1998). Moreover, because the aluminum industry is also energy–intensive, the plants are located near hydropower plants, and have enjoyed favorable long-term contracts with electricity suppliers and exemption from the electricity-tax (SSB, 2010; Vevatne et al., 2004).

In aluminum production, the Søderberg and Prebake anodes are used in the electrolysis process which consists of extracting the aluminum from the alumina driven by an electrical current (aluminum-production.com, 2019). The main differences between the Søderberg anodes and the Prebake anodes include: the latter is more efficient than the former, it uses less electricity (pot (cell) voltage), it needs less heat for anode baking, and the emissions per unit of aluminum produced are lower (Barber & Taberaux, 2014; 2014). These differences are summarized in Table 2.

Table 2: Comparisons of Prebake Technology to Søderberg Technology

Efficiency	Prebake > Søderberg
Pot (cell) Voltage	Prebake < Søderberg
Heat needed for the anode baking	Prebake < Søderberg
Emissions	Prebake < Søderberg
Production costs	Prebake < Søderberg
Heath damage	Prebake < Søderberg

Notes: Based on information from Kvande and Drabløs (2014), and Barber and Taberaux (2014)

The seven aluminum plants employed 4,691 people in 1996, more than 56% of the total number of employees in the metallurgical industry nationwide during same year (Godal, 1998). The metallurgical industry in Norway is not a labor-intensive industry. Yet, because the aluminum plants are mostly located in places with relatively low population density, they may be considered cornerstone companies in their associated municipalities as they are often crucial to employment in the local communities around the plants. In 1996, the employment in aluminum production ranged from 4.5% to about 19% of the total municipality's working age population (Godal, 1998, p. 10), as illustrated in Table 1. However, according to Godal (1998), these numbers may underestimate the importance of the business to the local community, as there are valuable spillover effects dependent upon the businesses' existence.

There are multiple by–products of the aluminum production process, which include per–fluorocarbon gases, in addition to carbon dioxide (CO_2) (Godal, 1998). The aluminum sector is one of the most carbon–intensive industries in Norway, but emissions decreased by 39% between 1990 to 2017 (Ministry of Trade, Industry, and Fisheries, 2000; Miljøstatus, 2018a). According to the Ministry of Trade, Industry and Fisheries (2000), most of the decrease in pollution these years was a result of technological and operational improvements that reduced the emissions of CO_2 related to production. Thus, the emission of CO_2 is proportional to production volume and new technology (Godal, 1998). In 2017, the metallurgical industry represented 43% of the climate gas emissions from the Norwegian manufacturing sector (Miljøstatus, 2018a).

2.3 Environmental Regulations of the Aluminum Industry

Since 1983, there seems to have been a relatively continuous pressure from the Norwegian Environment Agency (NEA) on the entire Norwegian manufacturing sector, which includes the aluminum industry.[8] The pressure has been related to improvements in production processes to reduce pollution. Meanwhile, the aluminum plants have taken the initiative to develop technologically, which has reduced the pollution from the plants (Miljøstatus, 2018a). The pro–manufacturing attitude in the government before 1997, as well as lobbying, resulted in an exemption of the aluminum industry from the CO_2–tax introduced in 1991 (Kasa, 2000). The exemption from the CO_2–tax was argued both against and in favor of, in different ways. Bye et al. (1999) estimate a relatively high marginal willingness to pay in the aluminum sector per tonne CO_2 emitted, while Groven et al. (1999) suggest that the manufacturing sector in Sogn og Fjordane was exempted from emission restrictions because the local aluminum industry had carried out measures to reduce emissions, satisfying the regulations of the Kyoto protocol. Moreover, it has been suggested that the means of taxation, like the CO_2–tax, would have had negative regional distribution effects (Reitan, 1998). Aluminum producing companies had, in 1996, few ways of substituting emissions related to production, and as an energy–intensive industry, the effect of the CO_2–tax on the price of electricity would threaten their operations (Bye et al., 1999).

In 1997, a voluntary agreement was signed between the aluminum producers and the MoCE. With the overall focus on domestic measures in the Norwegian political environment after 1981 and the replacement of government in 1997, uncertainty for the aluminum plant operators led them to come to a voluntary agreement as the best alternative to environmental measures in 1997. Thus, the motivation for accepting the agreement may have been to replace mandatory regulatory action plausibly imposed in the absence of a deal (Kolstad, 2011). Such voluntary agreement may, according to Vevatne et al. (2004, p. 1), a voluntary agreement may have had little effect on the ability to compete, and on the pollution level. In the 97–agreement, the emissions per tonne produced aluminum were to be reduced by 55% from 1990 to 2005 (Bruvoll & Bye, 2009). In 1998, the NEA took initiative to revise the old emission permits that had been given to the aluminum industry (NEA, 2000). This revision resulted in a stringency of the permits, to send a signal to the aluminum industry that immense improvements had to made, and that the

[8] NEA is a government agency under the MoCE whose primary tasks are to reduce greenhouse gas emissions, manage Norwegian nature, and prevent pollution (NEA, n.d.). It encompasses the former Norwegian Climate and Pollution Agency.

emissions from the Søderberg anodes would not be accepted in the future (Helgelendingen, 2000; Johansen, 2008). More precisely, each aluminum plant was left to decide whether to end production using the Søderberg anodes or to modernize the technology. The final compliance deadline for phasing out the Søderberg anodes was set to 2007 (Johansen, 2008).

Thus, As a result of the 97–agreement and the 2000–restriction (Gram, 2008; Hydro, 2007), five of the seven aluminum plants finally shut down their Søderberg anodes between 2003 and 2009 (Gram, 2008), and were left with Prebake technology that was considered better. Table 1 presents the different plants and their year of Søderberg–shutdown. As a result of the differences between the two technologies summarized in Table 2, the Søderberg cells were not only related to higher production costs, but also to a vulnerable point when required to reduce the emissions associated with production (Kvande & Drabløs, 2014).

3 Literature Review

Research has, since the late 1990's covered the debate regarding the cost–effectiveness of regulations designed to improve environmental quality. Positive externalities including health benefits of environmental regulation have been estimated by Shclenker and Walker (2011), and Chay and Greenstone (2003). These benefits, however, are not within the scope of the cost–effectiveness analysis for the regulated firm, as the firm's purpose is to minimize the costs associated with pollution reduction. The abatement costs related to environmental regulations have led manufacturers to argue that environmental regulations curse them with a competitive disadvantage in the global economy (Walker, 2011; Greenstone, 2002; Liu, Shadbegian, & Zhang, 2017). The costs related to these benefits might influence labor demand because a company facing increased marginal costs due to abatement requirements or other forms of regulatory compliance may be forced to adjust its input factors, such as labor (Kolstad, 2011). Yet, the consequences on labor demand are ambiguous according to the existing literature.

Previous empirical research documents a negative effect of environmental regulations on industrial location (Henderson, 1996), and of environmental regulations on employment level (Henderson, 1996; Greenstone, 2002). On the other hand, other research suggests that environmental regulations do not harm regulated firms or their workers in any significant way, and may even benefit them (Berman & Bui, 2001; Porter & van der Linde, 1995). Lastly, some claim that the labor reallocation costs related to regulation are more important than the effects on the level of employment (Walker, 2011; Walker, 2013).

As mentioned in Section 2, the transition from Prebake to Søderberg anodes allowed for increased efficiency and decreased pollution in the production of aluminum (Barber & Taberaux, 2014), and therefore improved environmental quality (Kolstad, 2011, p. 413-417). Labor, capital and natural resources are input factors used in the production of goods and services (Gärtner, 2016, p. 6). In aluminum production, the primary natural resource utilized is hydroelectricity. Hence, with more efficient anodes in the Prebake technology, there is need for less electricity to produce the same level of aluminum. This is in line with what Sharplin and Mabry (1986) define as technological advancement: "any change in a production process leading to higher standards of living (i.e. health benefits) through the increased output from the same amounts of resources or through the use of fewer resources to produce the same level of output" (Mabry & Sharplin, 1986, p. 2). Thus, observing this technological advancement, which is a form of technological change, we will present in the following section what researchers have found regarding the effects of environmental regulations on technological change, and of technological change on labor demand, in addition to presenting literature on the effects of environmental regulations on labor demand.

3.1 Environmental Regulation and Labor Demand

Previous literature on the effect of environmental regulations on labor demand shows that there is not a clear consensus on how environmental regulation affects labor market outcomes. Berman and Bui (2001) explore the effect on labor demand of the local air quality regulations introduced during 1979-92 in the Los Angeles Basin. To motivate their estimation equation, they design a theoretical model, "Labor Demand Under Environmental Regulation" (Berman & Bui, 2001, p. 274-276), that allows for regulations to act through both the output elasticity of labor demand and the marginal rates of technical substitution between abatement activity and labor. The output effect on labor demand is widely believed to be negative. However, the theoretical model does not give a clear prediction of the role of investments in abatement technology resulting from compliance, and the effect of these investments on marginals costs. This is further likely to affect employment through increased output. A change in the demand for environmental quality also often leads to a positive change in demand for abatement activities. However, the preceding sign of the effect on labor demand will, ultimately, depend on the type of abatement technology introduced, and whether it is a complement or substitute to labor. Thus, the true dominating effect of environmental regulation cannot be predicted from the model alone (Berman & Bui, 2001).

Berman and Bui (2001) find no evidence that local air quality regulation reduces employment, and argue that one of the reasons for this is because the regulated plants they analyze were intensive in capital, not in labor. By comparing the manufacturing industry in a locally–regulated region (treatment) to the same industry in a non–locally–regulated region (control) through first–differences estimation, they find that even though regulations impose large costs on the industry, they have limited effect on employment. They even find evidence of a slightly positive effect. Hence, the authors argue that environmental regulations can be labor–enhancing through abatement inputs which complement employment. Ferris, Shadbegian, and Wolverton (2014) investigate the environmental regulation of power plants in the US, and in particular, the employment effects from Phase 1 of a cap–and–trade program for SO_2–emissions implemented under the 1990 Clean Air Act Amendments (CAAAs). Their results, estimated using a differences–in–differences estimator, show little evidence of the Phase 1 power plants having a significant decrease in employment, relative to non–Phase 1 power plants. These results are consistent with the findings of Berman and Bui (2001).

However, there are scholars who come to contrary conclusions. Henderson (1996), Greenstone (2002), and Walker (2011) examine the inception of the Clean Air Act (CAA) of 1970 in the US, and the subsequent amendments in 1977 and 1990, while Liu, Shadbegian, and Zhang (2017) exploit stringency of wastewater discharge standards in China. They all find adverse effects of regulations on industry employment levels.

Recently regulated production is typically relocated to areas or sectors that have a record of staying in attainment, presumably to reduce regulatory scrutiny (Henderson, 1996; Greenstone, 2002; Walker, 2011). Greenstone (2002) applies a differences–in–differences model to examine the effect of a county's designation as out of attainment with the National Ambient Air Quality Standards for criteria of air pollutants on employment. He finds that nonattainment counties lost roughly 600,000 jobs relative to attainment counties over a 15–year time period. Walker (2011) also investigates the effect of the CAA on manufacturers, but looks at the amendment from 1990 as opposed to Greenstone (2002). He examines the changes in regulations following this amendment, which led to the Environmental Protection Agency's (EPA) adoption of new and more stringent pollution standards. He uses data on plant–level regulatory status linked with the Census Bureau Longitudinal Business Database, and decomposes net changes in employment due to new and more stringent pollution standards into job–flow components through a triple–differences approach. Walker concludes similarly to Greenstone (2002), finding that the strengthening of emission standards led to a persistent decline in employment in affected sectors, driven primarily by an increase in the plant–level job destruction rate. Liu, Shadbegian and Zhang (2017) support these findings through examining the impact on labor demand of a more stringent wastewater discharge standard imposed on all the textile printing and dyeing enterprises in a particular region of China. Through a differences–in–differences framework they find that the new environmental regulation reduced labor demand by seven percent.

More recent work by Walker (2013) has addressed the insufficient examination of the costs of labor market adjustment by looking at labor reallocation, and not only examining employment losses. He analyzes the effects from the CAAAs of 1990 also used in previous research (2011), and the impact on the transitional costs associated with a reallocation of workers from newly–regulated industries to other sectors of the economy. He finds that in the years after the introduction of the policy, the workers in newly–regulated plants experienced foregone earnings of about 20%, compared to their past earnings, driven by non–employment or lower earnings in future employment. Despite this, by highlighting the estimated health benefits of environmental regulations, Walker (2013) argues that the benefits far exceed the costs in total.

3.2 Environmental Regulations and Technological Change

The effects of environmental regulations on technological change have been investigated by several researchers (i.e., Ashford & Heaton, 1983; Milliman & Prince, 1989; Porter & van der Linde, 1995; Dupuy, 1997; Jaffe, Newell, & Stavins, 2002; Popp, Newell, & Jaffe, 2010), and their results suggest there is a causal relationship between policy instruments employed by governments and the technological responses of firms. Environmental policy has played a key role in inducing firms to adopt pollution–control technologies (Dupuy, 1997). Furthermore, Ashford & Heaton (1983) suggest that both the directly regulated industry and its related firms, particularly the suppliers of compliance technology, will create new activity patterns as a consequence of regulation. Such patterns would encompass new innovations as cleaner production methods, which makes it possible to achieve a specified level of reduced emission of pollutants in a more efficient way. Jaffe et al. (2002) refer to this process as the "induced innovation" hypothesis (Jaffe et al., 2002, p. 44-45).

The costs of environmental regulation and policy can include: decreased output of desired products, increased use of other variable input factors, purchase of specialized pollution–control equipment, or substitution of inferior or more expensive products or production methods (Popp et al., 2010). Innovation offsets from regulation have the potential to lower production costs, reduce the marginal cost of achieving a given unit of pollution reduction (Popp et al., 2010), and lead to a competitive advantage towards firms in foreign countries subject to similar regulations (Porter & van der Linde, 1995, p. 98). Popp et al. (2010) argue that in the absence of environmental policy, firms have little incentive to install cleaner technologies or to develop better environmental technologies for future use. The exception would be when the benefits of implementing a technological change are not only a pure public good, but a good for the user of the technology (Popp et al., 2010, p. 876-877).

3.3 Technological Change and Labor Demand

A growing body of literature argues that the reallocation of tasks between capital and labor has played a key role in reshaping the structure of labor demand in industrialized countries in the recent decades (Mabry & Sharplin, 1986; Manning, 2004; Acemoglu & Autor, 2011; D. H. Autor, 2013). Mabry et al. (1986) argue that technological advancement might lead to short–run effects on unemployment when there are imperfections in labor or product markets (Mabry & Sharplin, 1986). However, in competitive markets, technological advancement in one industry would release labor resources that could be allocated to other industries. Thus, a possible increase in unemployment due to technological advancement should be temporary (Mabry & Sharplin, 1986). Nonetheless, in the literature from the current century, a consensus of the effect of skill–biased technical change has formed among economists (Manning, 2004).

As an example, Autor et al. (2003) show that within industries, occupations and education groups, computerization is associated with reduced labor input of routine manual tasks and routine cognitive tasks, and increased labor input of non–routine cognitive tasks. Their findings are estimated through a model that demonstrates how an economy–wide decline in the price of computer capital affects task demand. They conclude that the effect from a technical change on labor demand depends on whether workers are direct substitutes or complements to the new technology. Substitutes suffer from the introduction of new technologies, while complements benefit from productivity increase. The theory is further expanded by Acemoglu and Autor (2011), who through their task–based framework based on the canonical model, demonstrate that technical change favoring one type of worker can reduce the real wages of another group (Acemoglu & Autor, 2011). Later, Autor (2013) shows "that many of the middle–skill jobs that persist in the future will combine routine technical tasks with the set of non–routine tasks in which workers hold a comparative advantage—interpersonal interaction, flexibility, adaptability and problem–solving" (2013, p. 196). Reviewed in conjunction with the literature presented in the previous subsection, the true consequences from a technological change originating of environmental regulation should be evaluated in relation to the tasks replaced by the pollution abatement technology.

3.4 Implications of the Study

This thesis aims to estimate the net effect of environmental regulation on outcomes of labor demand, as opposed to for instance Greenstone (2002), who estimates the gross effect. We estimate the net effect as we apply labor market outcomes at the municipality–level as proxies for the labor demand in the aluminum industry. Moreover, we use such data as proxy for the generated income in the municipality. It represent the first study that apply principally municipality–level data when estimating the effects on labor demand, and the special position of cornerstone businesses allows for us to create consistent estimates. To the best of our knowledge, the relationship between pollution, environmental regulation and labor demand, and the relationship between technological change and labor demand, has not yet been studied in a Norwegian, nor in a Nordic, context. Only Groven et al. (1999) have investigated barriers and possibilities for the development of climate policies with a basis in the local Norwegian level (Groven, Lundli, & Aall, 1999). The lack of research supports the notion that our contribution is necessary and relevant for policy–makers. Our work adds to the existing literature by investigating the effects of regulatory stringency in the aluminum industry in Norway. Walker (2011) emphasizes that much of the previous literature focused on earlier time horizons when pollution levels were much higher and technological constraints greater. Thus, the existing estimates of the effect of regulation on labor demand may no longer be applicable in today's economy. Our time period of interest is closer to the present. Therefore, this study may still be relevant in today's economy and highly relevant for current national and international policy debates regarding environmental protection.

4 Data Description

To investigate the effect of environmental regulation on labor demand, we have constructed a set of panel data that consists of municipality and labor characteristics. It is done by linking labor market variables on municipality–level with the environmental regulations and control variables. The primary source of data is the Norwegian Centre for Research Data (NSD), supplemented with data collected from Statistics Norway (SSB).[9,10] All data is collected on the municipality–level.

Inspired by the reviewed literature with a DiD approach (Beerli & Peri, 2015; Walker, 2013; Schlenker & Walker, 2011; Havnes & Mogstad, 2011), the sample consists of observations from five years before the cutoff. With the 97–agreement as our policy change of interest, the five years before the agreement, 1992 to 1996, has been contracted to assess the validity of the assumption of parallel pre–treatment trends in our DiD model. To adjust the sample to our IV model and to account for the last Søderberg shutdown in 2009, we include data up to 2010. Hence, we use data from 1992 to 2010 in our estimations. Furthermore, the sample is expanded to 2013, to make it possible to run a sensitivity check of the IV identification strategy in Section 6. Therefore, all variables are obtained from 1992 to 2013.

Additional to the population characteristics presented later in this section, we have collected data from SSB on whether the municipalities have a coastline or not. This dummy variable is time–invariant and is utilized to exclude municipalities without a coastline from our sample, which we get back to in Subsection 5.1. A list of the obtained population characteristics, outcome variables and control variables, and their descriptions, are presented in Table 8 in Appendix 8.

4.1 Data on Labor Demand

In our analysis, we explore movements in labor demand on municipality–level and assume they reflect the labor demand in the aluminum industry. The outcome variables of interest are: the unemployment rate, the employment rate, and the payroll tax paid to the municipality per person in the working age population. The industry is an important employer in the municipality, as explained in Section 2, thus, as a cornerstone business, we expect to see the changes in labor demand on a plant–level mirrored in the labor market outcomes on a municipality–level. Moreover, we presume there are synergistic

[9] Data is collected from NSD unless else is specified.
[10] NSD develops a municipality–database that has existed since the middle of the 70s (Nilsen, n.d.)

effects to the local labor market from movements in the aluminum industry. The most evident labor force indicators to study are employment and unemployment, as the labor force can be measured as the sum of these (Borjas, 2016). We expect an eventual change in employment to respond along three margins. First, employed workers may remain employed, either in the same job or another type of job. Second, employed workers may transition into or out of unemployment, claiming unemployment benefits or some other measure of unemployment. Third, employed workers may transition into, or out of, the labor force, a transition that is not likely to happen in the short–run but might be a consequence in the long–run. Holding municipality–level data as the lowest level of data available, two of our three main outcome variables of interest, the unemployment rate and the employment rate, serve as proxies for the number of employed persons at plant–level in the aluminum industry subjected to the environmental regulations. Moreover, we use payroll tax per person in the working age population as a proxy for generated income within the municipality.

According to SSB (2019a), a person who has no income-earning work, is actively seeking employment, and is available for work within two weeks, is considered unemployed. Also, an involuntary leave of absence is recognized as unemployment after a continuous duration of three months or more. In Norway, there are two different ways of measuring the unemployment rate (Bø & Næsheim, 2015). On one hand, the Labor Force Survey (LFS) takes into account all persons unemployed (SSB, 2019a), including those not registered with the Norwegian Labor and Welfare Administration (NAV). On the other hand, the statistics developed by NAV only measure persons registered as unemployed with NAV. The latter serves as a provider of information on how many individuals are searching for jobs through NAV, and on how many are receiving unemployment benefits. The statistics from the LFS are more comparable to international measures, and arguably more accurate (Bø & Næsheim, 2015). Moreover, utilizing the rate instead of the real number, enables comparison across municipalities in our analysis. Therefore, we use the unemployment rate calculated by LFS in our analysis. This number on the unemployment rate is available for the period 1993 to 2013, lacking 1992.

The data collected on employed persons includes the number of all employed persons in the age range 15 to 74.[11] These observations are based on information from LFS (SSB, 2019a). According to SSB (2019a), "Employed persons are persons aged 15-74 who performed work for pay or profit for at least one hour in the reference week, or who were

[11]The number excludes self–employed, and seamen. But seamen are only included in 1992. The age range was 16 to 74 from 1992 to 2005, this is accounted for in both Employed and Working Age Population.

temporarily absent from work because of illness, holidays etc."[12] To obtain comparable figures, the employment rate is calculated as the ratio of the number of persons employed to the working age population in each municipality. The employment rate can be used as a proxy to employed persons on plant–level. Because the number of unemployed is only available from NAV, we prefer the ratio to the working age population to calculate the number employed as a ratio of the labor force based on the number of employed from LFS plus the number of unemployed from NAV.

Our third outcome variable of interest, the payroll tax, is calculated as a share from employees' gross salary, and is paid by employers (Store Norske Leksikon, 2017). The payroll tax is paid to the municipality where the business is registered in the Central Coordinating Register for Legal Entities (Kjensteberg & Behringer, 2019). In Norway, the payroll tax is subject to regional differentiation as a mean to stimulate settling and employment in the districts, and the size is larger in centralized areas (Hervik & Rye, 2010). This implies that the municipalities of interest in our sample are divided into different zones of calculation. In 2007, Høyanger, Årdal, Sunndal and Kvinnherad (Husnes) were in the same zone experiencing a payroll tax of 10.6%, Karmøy and Farsund (Lista) in another experiencing a tax rate of 14.1%, and Vefsn (Mosjøen) alone was subject to a lower tax rate of 5.1%. The payroll tax rate is utilized as a proxy for generated income in the municipality, and is expected to reflect the condition of the labor market and labor demand. The measure is CPI–adjusted to 2018 Norwegian kroner (NOK) to enhance readability and allow for comparison over time. To facilitate comparison of the payroll tax across municipalities, we calculate the payroll tax per person by dividing the total payroll tax paid in each municipality by the number of persons in the working age population. As the payroll tax is paid as a percentage of each employee's salary, the payroll tax as a ratio of the working age population enables us to compare the relative change in the size of the total income at a municipality–level across municipalities. The variable is available for the entire period of interest, from 1992 to 2013, but there are missing values for some municipalities in the years 1992 to 1994.[13]

To use the number of employed persons and payroll tax as ratios to the working age population to enable comparison, we have defined a measure of the working age population. It takes into account all the persons in the working age population in the municipality, encompassing both those participating, or who could be participating, in the labor force according to their age. We consider all persons aged 15 to 74 as part of the working

[12]There is data available separately for employed men and women for the entire period 1992 to 2013. The two datasets are merged to get a measure of all the employed persons in each municipality.

[13]Vefsn is the only municipality of the treatment municipalities affected by this.

age population (SSB, 2019a).[14] The working age population variable is not equivalent to the labor force, because the labor force is the sum of employed plus unemployed persons (Borjas, 2016), while the working age population also includes those who are not in the labor force. An example of someone excluded from the labor force is someone who voluntarily does not participate, or who is not able to participate due to a high degree of disability. To apply this variable in our analysis and calculation of the employment rate and the payroll tax, we assume that the treatment is not correlated with the number of people who voluntarily do not participate in the labor market in the short–run.

4.2 Data on Control Variables

Several variables are collected to control for observable demographic and labor market characteristics, for the period from 1992 to 2013 on a municipality–level. Unemployment is concentrated among particular demographic groups and among workers in specific sectors of the economy (Borjas, 2016, p. 500-502). There are typically higher unemployment rates among less educated workers, young people, and in the manufacturing sector relative to the service sector. Traditionally, the unemployment rate among women has also been more elevated than that among men. Such characteristic differences can be the reasons for 'unemployment gaps', or employment differences, between the control and treatment groups in our DiD model, which gives us reason to control for these if there is a noticeable difference between the two groups in both the pre–treatment period and over time.

The population density denotes the number of inhabitants per square kilometer. By including this variable, we can compare the average population density in our treatment and control group, and assess whether the numbers are relatively similar. Hence, we can ensure that our decision to adjust the sample based on the average population amount, as done in Subsection 5.1.1, is coherent.

Data on the number of women living within each municipality in each time period has been divided by the population of the municipality, to create a variable for the share of women in the population. Similarly, an alternative to the share of women is included—the share of women employed. This is arguably more relevant to our study. It is likely that the share of women in the municipality's population has been relatively stable in Norway overall within the last decades, while the share of women employed may still vary between the treatment and control groups. We create a variable for the share of youth in the work force, created by dividing the number of persons aged 15 to 24 by the working age

[14]Before 2006, the working age population includes all persons aged 16 to 74. Note that this is not in line with the Employment definition of, e.g. OECD, where the working age population is defined as a person aged 15 to 64 (OECD, n.d.). We have accounted for this, referring to at in Table 8 in Appendix 8.

population.[15] We have also defined children as persons aged 0 to 15, elderly as persons aged 75 and above, and calculated the ratio of the number of children within the total population of the municipality. Including these five share variables allows us to control for differences between the municipalities in the demographics of the population.

Four variables representing the level of educational attainment in the municipalities are obtained. These reflect how many inhabitants over the age of 16 that have completed primary education, high school education, higher education up to four years, and higher education over four years as their highest level of education. For more comparable numbers, shares of the four variables are calculated by dividing them by the number of individuals comprising the adult population—the population over 16 years.

As mentioned in Subsection 2.2, the aluminum industry is capital intensive and is therefore likely to be sensitive to large fluctuations in the aluminum price caused by international competition (Vevatne et al., 2004; Godal, 1998). We are concerned that large fluctuations in the aluminum price could bias our estimates, because changes in the price could force the plants to adjust production levels. Thus, we control for the average annual aluminum price in our estimations. Including the aluminum price allows us to control for the influence of the fluctuations in price on net income. The aluminum spot price is collected from SSB (SSB, 2019b), and is included as a proxy for changes in the net income of the aluminum plants.[16]

[15]The unemployment rate is higher among people aged 15 to 24 years than among the working age population (SSB, n.d.). This could be related to the fact that they are students, and transition in and out of the labor market.

[16]The aluminum price from SSB is based on numbers from Quandl, which further is based on data from the London Metal Exchange and other commodities exchange.

5 Empirical Approach

The goal of this thesis is to identify the causal effect of environmental regulations on labor demand. In this section, we explain the identification strategies used to obtain this effect. We take advantage of a policy implementation, namely, the voluntary agreement made between the primary aluminum industry and the MoCE in 1997 (i.e., the 97–agreement), and account for the long compliance period of this agreement and an additional regulation taking effect in this period by dividing the post–treatment period in two. First, we apply a DiD approach, exploiting the 97–agreement as the start of Phase 1, and the increased stringency in emission permits (i.e., the 2000–restriction) as the start of Phase 2. Further, we apply IV estimation to examine whether there has been an effect on labor demand stemming from the technological advancement of changing from Søderberg to Prebake technology in five of the seven aluminum plants in Norway.

5.1 Differences–in–differences Approach

To answer our question of research, we turn to a generalized differences–in–differences (DiD) model. This approach enables us to explore the dynamic effects of the 97–agreement on local labor demand in the associated municipalities. We exploit the fact that the environmental regulations were only designated to the seven municipalities in Norway with a plant producing primary aluminum. This enables segregation of these seven aluminum–producing municipalities from the remaining municipalities in Norway, and comparison across levels, and time, with a suitable control group to establish a causal relationship. The available municipality–level data from 1992 to 2010, containing both the years before and after the establishment of the environmental agreement, allows us to make use of this particular empirical approach. DiD estimation is most appropriate when the treatment, which in our case is the 97–agreement, is random, or when observable characteristics can be controlled for (Ferris et al., 2014, p. 531).[17] Our DiD estimator might be biased without utilizing a suitable group of control variables.[18] We will come back to the criteria we have utilized to select our control variables in Section 5.1.5.

[17]A standard OLS estimation without an interaction term, and a treatment variable with a reduced sample including only the aluminum plants, would not have external validity as our sample would be too small to draw inference. If we included the entire sample used in DiD, we would treat all municipalities in the same manner, and would not be able to capture the changes which only happen in the aluminum municipalities. A Regression Discontinuity Design requires a running variable with a threshold, which is not present in our case of study, and this method is thereby not appropriate. Moreover, because we analyze several time periods, instead of applying a POLS, First–differences, Fixed Effects or Random Effects models, we have included fixed effects among our covariates.

[18]Rubin (2008) argues that we can approximate a randomized experiment by choosing a suitably matched control group to eliminate or reduce this bias.

Ideally, we would have exploited a random assignment of a dramatic policy implementation affecting all industries in Norway to prove the causal effect of an environmental regulation on labor demand. However, as such event has not occurred, a DiD strategy allows us to exploit a policy implementation directed to the aluminum industry. In order to calculate the true average treatment effect on the treated municipalities, we would need the average post–treatment outcome for the aluminum municipalities in the absence of the 97–agreement, $E[Y_0(1)|D = 1]$, subtracting it from the outcome for these municipalities when the 97–agreement occurs: $E[Y_1(1) - Y_0(1)|D = 1]$. Nevertheless, $E[Y_0(1)|D = 1]$, the true counterfactual, is impossible to observe, as this would have been an identical copy of the seven municipalities not treated with any policy. Therefore, we select a control group of municipalities representing this counterfactual as closely as possible, and proceed with the DiD estimator. This empirical strategy hinges on comparing the trends in outcomes of interest for the treated aluminum municipalities (i.e., the treatment group), with the respective outcomes of interest for a subset of comparison municipalities (i.e., the control group), before and after the environmental agreement. We argue that the primary aluminum plant operators would not have implemented means to reduce their emissions in the same way as they did after the agreement, if the agreement had not taken place. Thus, the 97–agreement can be argued to be an exogenous cutoff in our DiD analysis.

The compliance period of the 97–agreement is arguably long, because the involved parties were required to reduce their emissions first by 55% before 2005. With such a long compliance period, we cannot expect that the effect of environmental regulation on the local labor market will necessarily occur close in time to the onset. Preparatory processes and expectations in the markets might induce an impact before the specific changes in production are made. Furthermore, the penalties of not complying to the 97–agreement might not be sufficiently strong for the change to be implemented fast enough to associate it with a causal effect. Therefore, we will control for a gradual implementation of regulatory requirements in the aluminum industry. Moreover, with the stringency of the emission permits in 2000, the effect from the 97–agreement may overlap with the impact of the new restrictions taking effect in the middle of the compliance period.

In this thesis, we account for the long compliance period and a plausible gradual response by introducing two dummy variables that represent two periods. We divide our post–treatment period in two phases, one phase is the years between the establishment of the 97–agreement and 1999, and the other is the years between the 2000–restriction, and 2010. This approach is inspired by Havnes and Mogstad (2011) and Beerli and Peri (2015): researchers facing a similar issue of a twofold post–treatment period when applying

DiD.[19,20]

5.1.1 Treatment and Control Group

A primary aluminum plant is assumed to have an impact on the local community within the municipality, both in terms of employment and spill–over effects (Vevatne et al., 2004). Our treatment group consists of all Norwegian municipalities which include a plant. Thus, the treatment group includes the municipalities Årdal, Sunndal, Høyanger, Farsund, Kvinnherad, Vefsn, and Karmøy. All of these municipalities have a coastline, and most of them are located by fjords near hydro–power plants, due to the large amounts of electricity required to produce aluminum. In addition, they are all located in the southwestern part of Norway, including Nordland county, and none of them are considered large cities.

The structure of the Norwegian industry varies between the cities and smaller municipalities, depending on the degree to which an area is centralized (Ministry of Education and Research, 2018). More specifically, smaller areas in the districts are characterized by dependency on resource–based industries. Thus, since the municipalities in the treatment group are relatively small and remotely located, we have applied the following three criteria to achieve a set of control municipalities as similar as possible to the treated municipalities. First, the control municipalities must have a coastline. Second, they should be located in the southwestern part of Norway, including Nordland county. Third, we want to exclude big cities to avoid capturing characteristics related to cities. Consequently, to define our control group, we only keep the municipalities located in the counties Aust–Agder, Vest–Agder, Rogaland, Hordaland, Sogn og Fjordane, Møre og Romsdal, Sør–Trøndelag, Nord–Trøndelag, and Nordland. We exclude the cities Kristiansand, Sandnes, Stavanger, Bergen, and Trondheim, and we only include municipalities with a coastline.

From Figure 3, we see that the subset of possible control municipalities complying with the three criteria described above are small in terms of population, as most of them have an average number of inhabitants below 18,000. Figure 3 also depicts small differences in the average population size of the municipalities across the two groups, except for the municipalities with a very small average population in the subset of possible control municipalities.

[19]Havnes and Mogstad (2011) apply a phase–in dummy to account for the fact that the introduction of childcare coverage had a gradual effect on a certain cohort of treated children, in addition to a post–period dummy covering the cohorts of children that were affected with full force by the child care expansion.

[20]Beerli and Peri (2015) apply a dummy variable for the first phase of the gradual integration of the Swiss labor market into the European labor market, and a second dummy variable for the second phase and the post–period of the integration.

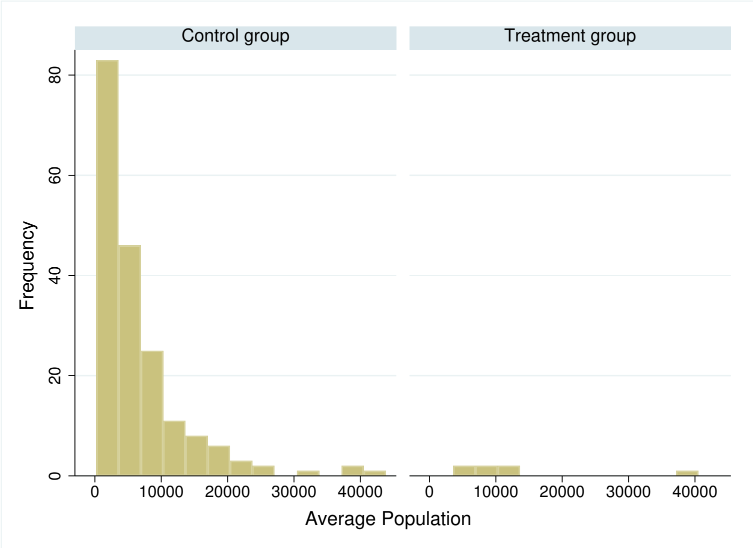

Figure 3: Frequency of Average Population across Sample Period in Control Group and Treatment Group

In addition to the three criteria presented, we want our treatment and control groups to be similar in terms of the average population across the sample period to avoid capturing characteristics of smaller municipalities. The lowest average population in our seven treatment municipalities is approximately 4,500. To exclude outliers from our dataset and to make the two groups as similar as possible, we therefore exclude the municipalities with an average population below 2,000. After excluding them, the control group amounts to 144 municipalities.

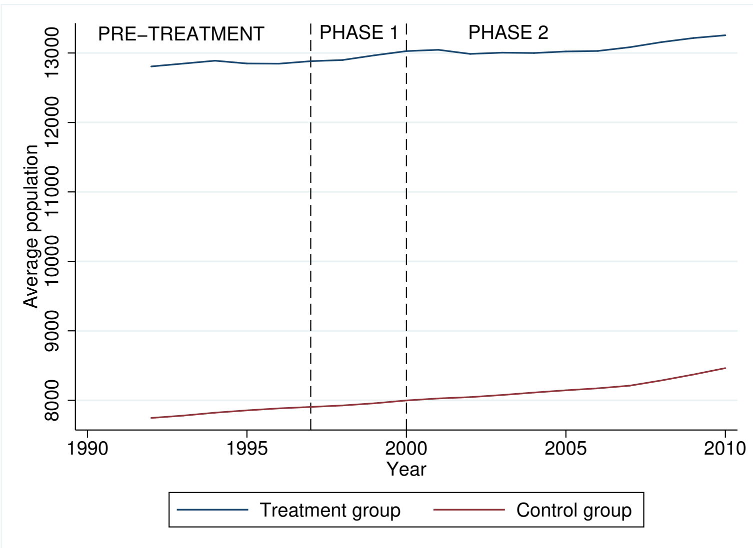

Figure 4: Trends in Average Population across Sample Period

The treatment and the control group are required to have parallel pre–treatment trends in outcome variables to ensure the internal validity of the DiD model (i.e., *parallel trends assumption*).[21] Figure 4 illustrates that the pre–treatment trends for the treatment group and the remaining observations comprising the control group, are reasonably similar, although the levels are different as the treatment group's average population is higher than that of the control group. Thus, we conclude that the selection process of a suitable control group has been successful, because all criteria are met. The treatment and the control municipalities are depicted in Figure 5.

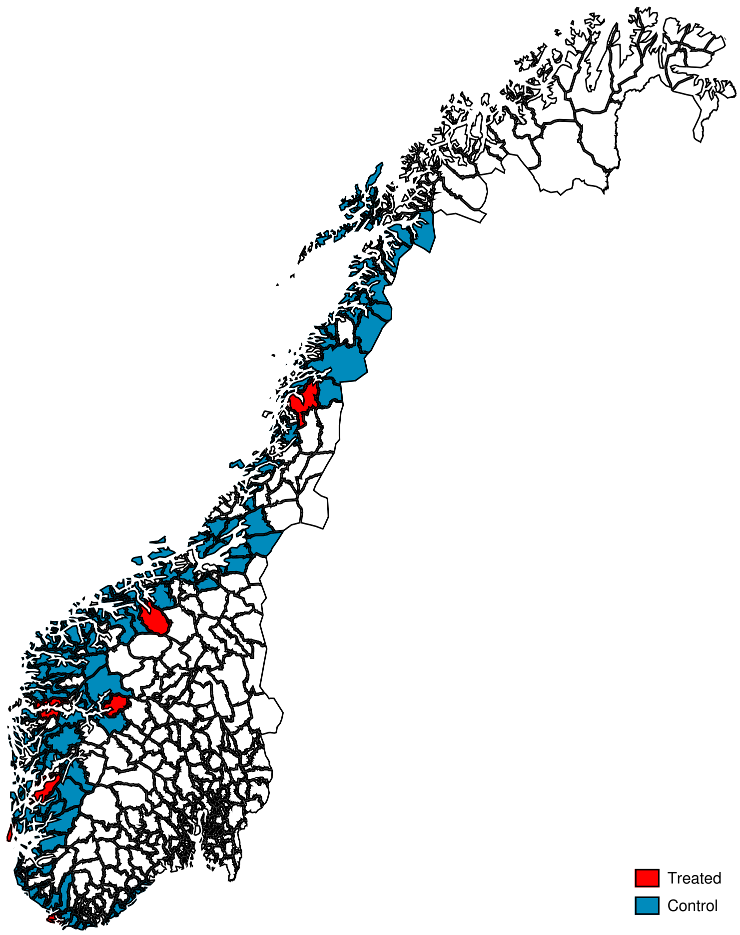

Figure 5: Map of Treatment and Control Municipalities

Notes: Treated municipalities are marked in red. These are municipalities with a primary aluminum plant. Control municipalities are marked in blue. The remaining municipalities are marked in white, and are excluded from the sample used for analysis.

[21]The discussion concerning this assumption is presented in Subsection 5.1.3.

5.1.2 Regression Model for the Impact of Environmental Regulations on Labor Demand

Our primary regression model, estimated by OLS over the sample of municipalities observed from 1992 to 2010, is defined as

$$\begin{aligned} y_{mt} = \alpha_m + \alpha_t + \gamma_1 Treat_m + \mu Phase1_t + \delta Phase2_t \\ + \gamma_2 (Treat_m \cdot Phase1_t) + \gamma_3 (Treat_m \cdot Phase2_t) + \epsilon_{mt} \end{aligned} \tag{1}$$

where y is the outcome of interest in municipality m and year t. The variable $Treat_m$ is a dummy variable equal to one for the treated municipalities, and zero otherwise, controlling for fixed differences between the treatment group and the control group (Angrist & Pischke, 2015). The two phase–variables control for the conditions changing over time. The variable $Phase1_t$ is an indicator dummy whose value is one in the years 1997-1999, corresponding to the first phase of the environmental regulations, and zero otherwise. $Phase2_t$ is a similar indicator equal to one in the years 2000-2010, encompassing the second phase, and zero otherwise. We include the interaction terms $(Treat_m \cdot Phase1_t)$ and $(Treat_m \cdot Phase2_t)$ to uncover the true causal effect of the environmental regulations. Our coefficients of interest are γ_2 and γ_3, and indicate the average causal effect on the outcome variables of interest in the aluminum municipalities in the respective time periods. The error term, ϵ_{mt}, reflects the idiosyncratic variation in potential outcomes that varies across time and municipalities.

If there is a correlation between the treatment variables, as well as between the outcome variables, and unobserved time and individual invariant characteristics, the estimates will suffer from unobserved–but–fixed omitted variables bias (OVB). This will in turn overestimate, or underestimate, the causal effect of an environmental regulation on the outcomes of interest. To correct for this bias, we include time and municipality fixed effects (Angrist & Pischke, 2009). The unobserved individual effects are coefficients on municipality dummies, represented through the municipality fixed effects coefficient α_m. This coefficient absorbs all time–invariant characteristics of each municipality, such as the physical characteristics. Other factors may not be exactly constant, but roughly constant over 15-20 years, such as initial sector specialization or population density. The year fixed effects are coefficients on year dummies represented by α_t, which absorbs yearly economic and demographic variation common for the whole sample. An example of such yearly changes are business cycles: for instance, the years most affected by the financial crisis in 2007-2008.

We aim to capture the causal effect on the local labor demand in the associated municipalities, which makes employment, unemployment or hidden employment interesting outcome variables for our analysis (Borjas, 2016). We include the unemployment rate as an outcome of interest, and the employment rate as a complement to this. Both measures are included to verify whether there are corresponding effects in each of the rates from the environmental agreement. As the plants are cornerstone companies in their respective municipalities, changes in the unemployment rate and employment rate at municipality–level are expected to reflect potential movements in the aluminum plants. Moreover, a third measure that tells us something about the labor demand is the payroll tax, calculated as a portion of employees' gross salaries which is paid by employers (Store Norske Leksikon, 2017). Thus, payroll tax can arguably be a proxy for the size of the income of the municipalities.

5.1.3 The Assumptions of Differences–in–Differences

To draw inference with the DiD approach, the following identifying assumption of parallel trends must hold. The *parallel trends assumption* entails that the outcomes of interest in the treatment and control groups should have evolved with the same pattern in the absence of the environmental regulation. This implies that the they must have followed the same trend before the introduction of the regulation. If they meet this criteria, the DiD approach allows us to control for unobserved differences between treatment and control municipalities, and a divergence from the trend in the treated sample in the post–treatment period can signal a treatment effect. Given parallel trends, the average treatment effect on the treated is identified as:

$$\begin{aligned} E[Y_1(1) - Y_0(1)|D=1] = \{&E[Y(1)|D=1] - E[Y(1)|D=0]\} \\ -\{&E[Y(0)|D=1] - E[Y(0)|D=0]\}. \end{aligned} \tag{2}$$

Equation 2 depicts the difference between the outcome of the treatment group following treatment, and its potential outcome in absence of the treatment on the left–hand side. This equals the difference between the outcome of the treatment group and the outcome of the control group posterior to treatment, minus the difference between the outcome of the treatment group and the outcome of the control group prior to treatment.

In order to evaluate which of our chosen outcomes of interest comply with the parallel trends assumption, we have analyzed a series of line plots, one for each of the relevant outcome variables of interest. In Appendix A.1, we have plotted the average development over time for the treatment and control group in the unemployment rate, employment rate and payroll tax per person in the working age population. We investigate how parallel the trends in the two groups are in the pre–treatment period. The fact that our sample includes many municipalities and years allows us to relax the parallel trends assumption, that is, to introduce a degree of nonparallel evolution in outcomes between municipalities in the absence of a treatment effect (Angrist & Pischke, 2015, p. 184-186). Evaluating each of the line plots, we find that all the chosen outcome variables of interest seem to have sufficiently parallel pre–treatment trends in line with the key assumption, yet, they are not perfectly parallel. The trends in unemployment rate in particular, move in opposite directions up until 1995, but from that point onwards, they move very similarly up until 1997. Thus, due to not perfectly parallel trends for any of the variables of interest, it will be especially important to control for pre–existing economic characteristics of the municipalities through control variables.

Two requirements that accompany the parallel trends assumption, are the *no–anticipation assumption* and the fact that the determinants of environmental regulation should not be systematically correlated with potential labor market outcomes. The latter requirement is arguably not an issue in this analysis, as the only motivation for introducing environmental regulation should be to increase environmental quality, and not to improve labor market development. If the aluminum plants anticipated the 97–agreement, or the subsequent 2000–restriction, this hinders conclusion of causality as it would violate the no–anticipation assumption.

In order to verify the fulfillment of the parallel trends assumption and the no–anticipation assumption, and hence, to increase confidence in our identification strategy, we run the following sensitivity checks in Subsection 6.1.1: We perform tests of the parallel–trends assumption by looking at the differential evolution of our outcomes of interest in the pre–treatment period, investigate the movements in the working age population, we run a placebo test, and extract treatment municipalities singly from the sample in order to test the importance of each municipality for our estimated results.

5.1.4 Standard Error Issues

A pillar of traditional cross–sectional inference is the assumption that observations are independent, both over time and within groups (Angrist & Pischke, 2009, p. 293). If observations are serially correlated, or dependent within groups, it can affect statistical inference severely. Average employment and unemployment rates are correlated over time, which causes serial correlation in our standard errors (Angrist & Pischke, 2009, p. 294). According to Bertrand et al. (2004), this issue is especially present in a DiD context. Another source of serial correlation is the presence of area–year random effects. For example, a tax reform can be a shock that affects only a certain county. Such random effects result in a violation of the parallel trends assumption. Furthermore, there might be correlation across sectors within the same county. This is referred to as the 'clustering problem'. If we fail to correct for this, the standard errors of our estimated coefficients will most likely understate their standard deviation (Bertrand, Duflo, & Mullainathan, 2004, p. 251). Bertrand et al. (2004, p. 251) propose clustering of standard errors as a solution to adjust standard errors for the correlation within groups evolving from such area–year shocks, and also to correct for correlation within the area–year shocks themselves. However, by collapsing the time series dimension into three periods, pre–treatment, Phase 1 and Phase 2, we reduce this problem considerably (Bertrand et al., 2004).

In our analysis, it could have been relevant to cluster standard errors at a county level, to account for possible correlation between regulatory status for industries within the same county.[22] This strategy would have allowed for shocks common to each county, and for dependency over time. However, most features related to labor demand and the plants are set on the municipality–level, making the municipality–level relevant for clustering. According to Wooldridge et al. (2017), there is a requirement of heterogeneity in the treatment effects for the clustering adjustment to be necessary if fixed effects are included at the level of relevant clusters (Abadie, Athey, Imbens, & Wooldrigde, 2017, p. 3). Hence, by correcting for heteroskedasticity with robust standard errors, we already handle the clustering in our standard errors. Thus, we do not need to adjust standard errors for clustering (Abadie et al., 2017).

5.1.5 Choice of Control Variables in the DiD Estimation

Already existing economic characteristics that could cause treated and control municipalities to differ systematically are essential to account for in our model, as they might over, or underestimate, the causal effect of interest like mentioned in Section 4. The evolution in the unemployment rate could, for instance, be driven by the share of young workers in the municipalities, resulting in certain systematic differences between the treatment and comparison municipalities that are important to capture. This plausible matter of endogeneity due to OVB will lead to a lack of confidence on whether changes in the outcome of interest, y, are driven by the explanatory variables, or by variation in unobservable or observable characteristics absorbed by the error term, ϵ_{mt}. To control for this, we extend Equation 1 to include a vector of covariates $\mathbf{X}_{mt}$, encompassing our set of control variables.

[22] Høyanger and Årdal are both located in Sogn og Fjordane county.

By the logic of OVB, several variables could be used to correct for the problem of omitted variables affecting the outcomes of interest. However, these control variables should also be relevant to our model. Bad controls can create selection bias (Angrist & Pischke, 2015, p. 215), and thus, the control variables must be chosen carefully. Angrist and Pischke (2015) emphasize that good control variables should not be possible outcome variables in our model, to ensure that they are unaffected by the treatment variable. Moreover, controls measured before treatment are generally good controls, because they cannot be changed by the treatment (Angrist & Pischke, 2015). In Subsection 5.2, relevant variables for our topic are presented, the source they are collected from, and conclusions regarding the specific controls we include are drawn.

To test for inadequate controls, we will also run the model without the control variables. If there is a significant change between the baseline Equation 1 with and without the set of controls, it implies that the variation in the outcome of interest is mainly explained by the controls, and not by the treatment variable of interest relied on by the identification strategy. In such a case, there may be reason to question whether the treatment variable has a direct effect on our chosen control variables. This will further be discussed in Section 7.

Contemporaneous Events

The causal interpretation of γ_2 and γ_3 as regulation effects relies on the identifying assumption that there are no omitted time–varying effects with different impacts on the outcomes of interest in the treatment and control groups. Thus, the DiD estimates might be biased if there are contemporaneous events or policy changes affecting municipalities in different ways. For instance, large tax reforms could bias our estimates on payroll tax. There have been two large tax reforms in our sample period, observed in 1992 and 2006 (Store Norske Leksikon, 2015). However, if these reforms were implemented in all municipalities simultaneously, and imposed equal changes on the respective taxes paid across the country, the effect will be absorbed in the year fixed effects. Across our sample, there seem to be an even distribution across payroll tax zones, and as both treatment and control municipalities are distributed evenly across the different zones, we do not recognize this as a problem. Furthermore, we have reason to believe that changes in payroll tax affect all zones equally. Another point to mention is that to the best of our knowledge, there has not been the establishment of any new major employers in any of the municipalities surrounding the treated municipalities in our sample period of interest.

5.2 Descriptive Statistics

Table 3 displays the means of our outcome variables of interest, as well as demographic and labor market characteristics. The DiD estimators are based on the presumption of time–invariant, or group–invariant, omitted variables (Angrist & Pischke, 2009, p. 243). We use the group mean of the treated and control municipalities to check whether our treatment and control group indeed look similar. As demographic characteristics should be unchanged over time, we expect to see only small differences between the control and treatment municipalities (Angrist & Krueger, 1999). Yet, to verify that there are not changes over time in the differences between the treatment and control group, we include the differences over time in the control variables. However, we will apply the characteristics that prove to be somewhat different between the two groups as our control variables in our equation of interest in the DiD approach, Equation 1. Furthermore, they will be extended as controls to our estimations in the IV approach. The purpose of the chosen control variables is to capture the differences in the labor market between the treatment and control municipalities, and in this way, to facilitate consistent estimates of the regulation effects in both estimation methods.

The outcome and control variable means displayed in Column (1) in Table 3 are for the treatment municipalities in the pre–treatment period. These means subtracted from the means of the control municipalities in the same period are depicted in Column (2). The difference between the two groups in Phase 1 and Phase 2 are further depicted in Column (3) and Column (4). As reported in the table, there is a significantly lower number of observations in the treatment group versus the control group because we have few treated municipalities.

From Panel A in Column (2), it is evident that there are small differences between the outcomes of labor demand for the treatment and control municipalities in the pre–treatment period. For the unemployment rate and the payroll tax per person, the values in the control municipalities diverge slightly in the pre–treatment period. Reverse movements are reported in the differences in the payroll tax per person and the employment rate from Phase 1 to Phase 2, and suggest substantial treatment effects in our DiD estimations. When graphing the outcomes of interest for the two groups over our sample period (i.e., Figures 6–8 in Appendix A), the movements correspond well with what we observe in Panel A.

Although we control for unobserved differences between the groups before and after the 97–agreement in applying the DiD approach, this strategy might be called into question if there are substantial changes over time in the differences of observable characteristics between the treatment and control municipalities. This does not seem to not be an issue of importance in our analysis when we observe the changes over time in Panel B for the different control variables. However, we observe that the share of women employed, elementary school and high school as the highest level of educational attainment are somewhat different in the treatment group compared to the control group when examining the differences in means more thoroughly. Consequently, we include these three variables as controls in our estimations of Equation 1 in the DiD approach, and of Equations 3–5 in the IV approach, beside the aluminum price. These are good control variables in line with the description of good controls as we trust that these are relatively stable over time, and that possible changes in the former two will take time to observe (Angrist & Pischke, 2015). Hence, we do not expect them to change as a result of the environmental regulation, and then are relevant factors that potentially could drive labor demand.

5.2.1 Alternative Specifications

We perform a placebo–test in Subsection 6.1.1 to investigate the assumption of parallel pre–treatment trends, pretending that the environmental regulations took place in the pre–treatment period. If there are significant effects from the placebo regulations, this would suggest that our estimated treatment effects are absorbing differential time trends, rather than a true policy impact.

Table 3: Descriptive Statistics: Outcome and Control Variables

	(1) Treatment pre–1997	Differences Control-Treatment (2) pre–1997	(3) Phase1	(4) Phase2
A. Outcome Variables				
Unemployment Rate	0.0493	-0.0060	-0.0022	-0.0018
Employment Rate	0.5800	-0.0474	-0.0338	0.0082
Payroll Tax	18285.48	-5185	-4610.12	1895.38
Observations Treated		26	21	77
Observations Comparison		550	417	1572
B. Control Variables				
Population Density	33.65	3.42	3.97	5.055
Share of Women	0.4950	0.0000	0.0010	0.0005
Share of Women Employed	0.4380	0.0248	0.0180	-0.01162
Share of Youth	0.1896	0.0073	0.0141	0.0087
Share of Children	0.2128	0.0104	0.0100	0.0063
Share of Elder	0.0735	0.0082	0.004	0.0004
Elementary School	0.3838	0.0324	0.0283	0.0222
High School	0.4837	-0.0308	-0.0300	-0.0299
Higher Education Long	0.0179	-0.0010	-0.0010	0.0009
Higher Education Short	0.1062	0.0030	0.0050	0.0072
Observations Treated		35	21	77
Observations Comparison		720	432	1584

Notes: Pre–treatment period is 1992 to 1996, Phase 1 period is 1997 to 1999, and Phase 2 period is 2000 to 2010. Column (1) depicts the characteristics of the treatment group in the pre–treatment period. Column (2) depicts the characteristics of the treatment group subtracted from the characteristics of the control group in the pre–treatment period. Column (3) depicts the characteristics of the treatment group in Phase 1 subtracted from the characteristics of the control group in the same period. Column (4) depicts the characteristics of the treatment group in Phase 2 subtracted from the characteristics of the control group in the same period. Outcome variables are defined in Section 4.1 and Control Variables in Section 4.2. All variables are yielded as shares, except for the payroll tax and the population density.

5.3 Instrumental Variables Approach

In order to investigate the effect of technological advancement on labor demand, we apply an instrumental variables (IV) approach on the same sample used in the DiD approach. In this study, the technological advancement represents the replacement of the Søderberg anodes with the Prebake anodes in five of seven of the aluminum plants. According to Berman and Bui (2001), their model on labor demand under environmental regulation, described in Section 3, would ideally have been applied with regulatory change variables as instruments for the abatement activities.[23] The principle abatement activity conducted by the aluminum plants, the Søderberg–shutdown, is expected to have affected labor demand, and represents a technological advancement implemented in five of the seven plants. In the following section, we apply a regulatory change variable as an instrument for this abatement activity.[24]

As described in Section 2, the 2000–restriction subsequent to the 97–agreement, led to the final shutdown of the Søderberg anodes between 2003 and 2009 (Johansen, 2008; Helgelendingen, 2000; Gram, 2008). Consequently, the 97–agreement did not enforce the shutdown. Thus, there is an endogenous technology advancement likely to affect our outcomes of interest. We initiate our analysis by using the increased emission restrictions to isolate exogenous variation in the shutdown of Søderberg anodes. Specifically, we instrument the years t after the increased stringency of emission permits on the years t after the endogenous technology advancement, the Søderberg–shutdown, for each municipality m. This is conducted by estimating the first–stage equation

$$Soderberg_{mt} = \alpha + \gamma Post2000_t + \beta_m t + \epsilon_{mt}, \tag{3}$$

where the dependent variable, $Soderberg_{mt}$, is equal to one for the years t after the Søderberg anodes are shut down in the respective municipality m, and zero otherwise. The independent variable $Post2000_t$ is the instrument, equal to one for each aluminum municipality in each year t after and including 2000, and zero otherwise. In addition, we include a municipality–specific linear trend coefficient, β_m multiplying a time trend variable, t. This is a generated number unique to each observation, specific for both year and municipality, which allows each municipality to follow different trends in a limited, but potentially revealing way. The IV strategy hinges on the assumption that the instrument

[23] Berman and Bui (2001, p. 275) conclude that it was "too ambitious a demand to make of [their] data".

[24] The rollout strategy is not an alternative approach to our study: Estimating the effect of the rollout of the Søderberg–shutdowns alone on outcomes, will give too little variation in our estimation as there are few observations experiencing a shutdown.

is relevant, often referred to as the *relevance condition*, entailing that the instrument must have some power to influence the treatment variable. Its relevance is verified through the first–stage regression, calling for a statistically significant first–stage estimate (Angrist & Pischke, 2015, p. 145). If the instrument is not highly correlated with the treatment variable, it implies that our first–stage estimate is insignificant and the instrument weak. Thus, it cannot be used to interpret causality as it will be biased and very sensitive to a small sample size. Through the first–stage we obtain the two–stage least square (2SLS) estimator, which is a new fitted value of the explanatory variable. Further, we proceed with the relationship between our outcomes of labor demand and the Søderberg shutdown dates, the second–stage, given by the following equation:

$$y_{mt} = \alpha + \mu \hat{Soderberg}_{mt} + \beta_m t + \epsilon_{mt}. \tag{4}$$

The predicted variable $\hat{Soderberg}$ in Equation 4 encompasses the instrument from Equation 3, and the causal estimator of interest is μ.[25] In Section 6.2.2, we show the second–stage estimates and the OLS estimates, to compare the effect of the treatment variable instrumented and not instrumented, respectively.

The instrumental variables approach outlined above is based on the assumption that the only way the 2000–restriction has affected labor demand is through the Søderberg–shutdown (i.e., through the first–stage channel). This assumption, the *exclusion restriction*, is violated if the 2000–restriction affects our outcomes of interest through other channels. We cannot guarantee that that there are no other sources for the 2000–restriction to have affected labor demand than the Søderberg–shutdowns, yet, to the best of our knowledge, this assumption holds.

Another requirement for the IV approach, is that the 2000–restriction has an indirect effect on labor demand. The regression of outcomes of labor demand on the 2000–restriction is displayed through the reduced–form equation

$$y_{mt} = \alpha + \beta Post2000_t + \beta_m t + \epsilon_{mt}. \tag{5}$$

If there is no statistically significant effect of the increased stringency on our outcomes of interest, this implies that the causal relationship of interest is weak or nonexistent (Angrist & Pischke, 2015, p. 146). The coefficient of interest in the second–stage Equation 4 amounts to the ratio of reduced–form to first–stage parameters, and the reduced–form

[25] We run our IV estimations using STATA's built–in command to avoid the problem of incorrect standard errors from manual 2SLS.

equation is necessary to test the exclusion restriction. If the reduced–form estimate is statistically significant with no evidence of a corresponding first–stage, it suggests that there is a link between the 2000–restriction and the labor demand other than the shutdown of the Søderberg anodes. In such an event, the exclusion restriction would be violated (Angrist & Pischke, 2015).

Furthermore, the timing and introduction of the increased stringency in emission permits must be uncorrelated with the determinants of changes in labor demand, implying that the instrument would not be a good control variable (Wooldridge, 2016, p. 761). This assumption is referred to as the *independence assumption.* We argue that the timing of the introduction of the 2000–restriction is likely to be exogenous with respect to developments in labor demand. Moreover, the 2000–restriction is likely to only have been implemented in response to an exigency for improved environmental quality, and not as a mean to affect labor market development. Thus, we conclude that the instrument is as good as randomly distributed across our observed variables of interest.

5.3.1 Choice of Control Variables in the IV Estimation

In the IV approach, one must handle the matter of OVB with increased diligence, as the concern over OVB arises in a new form when applying an instrument for the endogenous treatment variable. The IV estimation is biased if an omitted explanatory variable is correlated either with the included explanatory variables, or with the instrumental variable (Murray, 2006, p. 118). Taking this into account, we control for share of women in the employed force, share of adult population with elementary school at their highest educational attainment level, share of adult population with high school as their highest educational attainment level, and the aluminum price. We add X_{mt}, including these municipality–level controls, to our first–stage, second–stage and reduced–form equations.

Due to few municipalities exposed to the Søderberg–shutdown, we may lack variation in our instrumented variable, $\hat{Soderberg}$, to explain the variation in labor demand. Thus, there is a risk that the controls will absorb most of the variation. To lower their variation, we could have included our controls as dummies, and not as percentages. However, as both the instrument and the instrumented variable are dummies, we would get a perfect collinearity issue and some of the variables would be omitted. This is also the reason a municipality–specific linear trend replaces the year and municipality fixed effects used in our DiD model.

5.3.2 Alternative Specifications

As a robustness check, we perform a test of endogeneity, and expand the dataset to 2013 including available observations. With one of the Søderberg–shutdowns taking place in 2009, which is close to the last observation year in our main dataset, we want to check if an expansion of the sample changes the estimates. These checks are presented in Subsection 6.2.

6 Empirical Analysis

The theoretical prediction by Berman and Bui (2001) implies that the effect from environmental regulations on labor demand is ambiguous, a conclusion which is further supported by Ferris et al.(2014). On the other hand, Henderson (1996), Greenstone (2002) Walker (2011), and Liu et al. (2017) find adverse effects of regulations on industry employment level. In this section, we investigate whether our estimations of the effect of environmental regulation on outcomes of labor demand are in line with any of these predictions. Proceeding to the IV estimation, we investigate if the implementation of a technological advancement as a consequence of stringency in emission permits substantiates the findings of our DiD estimations. In the following, we first present the main findings of the DiD estimations together with a sensitivity analysis. We then present the main findings of the IV estimations with a separate sensitivity analysis. In the sensitivity analysis we test the identification assumptions and run robustness test to identify possible weaknesses of our approach.

6.1 Main Results Differences–in–Differences

Table 4 presents our main results based on Equation 1. Included in the baseline specifications in column (2) are the control variables for municipality and year fixed effects. The causal effects of interest, represented in the coefficients of the interaction terms, respectively denoted $Treat \cdot Phase1$ and $Treat \cdot Phase2$, are compared with the pre–treatment mean of the respective outcome variable for the treatment group, depicted in column (1). The size of the estimated effects on the employment rate, the unemployment rate, and the payroll tax per person do not vary much across specifications.[26] Therefore, the estimates including the specifications (i.e., the control variables) are preferred, and we will interpret these in this subsection.

Panel A in Table 4 shows the estimations of the effects of the 97–agreement and the 2000–restriction on the unemployment rate. In the years 1997 to 1999, there are no substantial effects on the rate as the numbers are not significantly different from zero at conventional levels.[27] However, in the period 2000 to 2010, the effects are negative and statistically significant. The effect of the 2000–restriction is associated with a 0.007 percentage point decrease in the unemployment rate. This represents a decrease equivalent

[26] We have run an F–test of the chosen control variables and reject the null hypothesis of joint insignificance at a one percent level.

[27] To reject a null hypothesis which specifies that the coefficient of a variable is equal to zero, we use a conventional level of significance ($p < 0.05$, or $p < 0.01$).

to approximately 15.6% relative to the pre–treatment mean.

The results of the employment rate are presented in Panel B in Table 4. They suggest that there is a statistically significant effect from the 97–agreement on the employment rate in the years from 1997 to 1999, and also in the period after and including 2000. The effect from the 97–agreement on the employment rate is associated with a 0.015 percentage point increase in the first phase, and an approximately 0.01 percentage point decrease in the employment rate subsequent to the 2000–restriction. Relative to the pre–treatment mean that is equal to about a 2.7% increase and a 1.7% decrease, respectively. These findings indicate that more people are employed as a result of the agreement in 1997, but that people lose their jobs as a consequence of the 2000–restriction. Based on these results, we investigate whether the estimates may be biased due to changes in the working age population as a ratio to the total population in Section 6.1.1.

For the payroll tax per person in the working age population, we find that it increases significantly at a ten percent level in Phase 1, while it decreases significantly at a one percent level in Phase 2. The effect of the environmental regulations on the payroll tax is associated with an increase of 1,970.5 NOK per person in the working age population after the 97–agreement was signed, and a reduction of 2,209.8 NOK per person subsequent to the 2000–restriction. These effects are equal to a 10.5% increase and a 11.8% decrease in the payroll tax per person, respectively, relative to the pre–treatment mean.

Table 4: DiD Estimates for Outcomes of Labor Demand

	(1) Mean pre–treatment	(2) Baseline results	(3) Including controls
A. Dependent variable: Unemployment Rate			
	0.0458		
Treat · Phase1		-0.00131	-0.00129
		(0.00238)	(0.00236)
Treat · Phase2		-0.00706***	-0.00713***
		(0.00212)	(0.00210)
Observations	28	2718	2718
Adjusted R^2		0.746	0.750
B. Dependent variable: Employment Rate			
	0.573		
Treat · Phase1		0.0174***	0.0152***
		(0.00521)	(0.00477)
Treat · Phase2		-0.0121**	-0.00990**
		(0.00564)	(0.00504)
Observations	35	3020	3020
Adjusted R^2		0.936	0.945
C. Dependent variable: Payroll Tax			
	18730.3		
Treat · Phase1		1781.1*	1970.5*
		(1058.4)	(1083.8)
Treat · Phase2		-1857.8**	-2209.8***
		(744.0)	(772.8)
Observations	32	2805	2805
Adjusted R^2		0.889	0.899
Fixed Effects		Yes	Yes
Control Variables		No	Yes

Notes: Estimations are based on OLS of Equation 1 presented in Section 5. The outcome variables are defined in Section 4.1. There are four control variables: Share of women in the workforce, share of persons with high school as their highest educational attainment level, share of persons with elementary school as their highest educational attainment level, and the aluminum price. The sample consists of 151 municipalities over the sample period 1992-2010. Phase 1 is the years 1997-1999, and Phase 2 is the years 2000-2010. Column (1) depicts the mean values in the pre–treatment period in the treatment group. Column (2) shows the baseline results from equation 1, including year and municipality fixed effects. Column (3) depicts the estimates from equation 1, including year and municipality fixed effects, as well as control variables. Standard errors are in parentheses and robust to heteroskedasticity.

* p<0.10, ** p<0.05, *** p<0.01

6.1.1 Sensitivity Analysis of the DiD Estimation

In this subsection, we check the robustness of our findings by running a collection of sensitivity checks to identify violations of the identifying assumptions, and potential weaknesses of our DiD model.

Anticipation Effects

A possible pitfall in our analysis is the presence of anticipation effects implying changes in outcomes before the adoption of the environmental agreement in 1997. A reason for such anticipation effects could be the CO_2–tax introduced in 1991, might causing the operators to prepare for the introduction of alternative environmental policies. The conventional diagnosis is that this makes the treatment endogenous and violates the no–anticipation assumption, which may cause biased estimates. In order to test this assumption, we estimate the differential evolution of our outcome variables of interests for the treatment and control groups, and check whether the difference is statistically significant in the pre–treatment period.

Figures 9, 10, and 11 in Appendix A.3, plot the coefficients of an interaction term between a year dummy variable and our treatment variable $Treat_m$ from Equation 1 at the municipality–level, and the 95% confidence interval. The plots of the employment rate and payroll tax per person show no differential trends before 1997, although the unemployment rate may have a tendency of an increasing differential evolution. Yet, the coefficients of the interaction terms are not significantly different from zero for any of the outcome variables of interest after 1997. Thus, there seem to be no significant anticipation effects in the outcome variables of interest. In addition, we confirm that the assumption of common pre–treatment trends holds.

Time Trends

Although we have confirmed that there are parallel pre–treatment trends in the check for anticipation effects, as well as in Subsection 5.1.3, underlying trends coinciding with environmental regulations might be a problem in our analysis. Such trends might be captured rather than a true policy impact, which potentially can bias our treatment effects. We investigate this issue by running a placebo test, simulating that the environmental regulation occurred in the pre–treatment period. Specifically, we add interaction terms between the treatment municipalities, as well as additional dummy variables for placebo treatment years, equaling one for the year 1995 and one for the year 1996, to Equation 1. If the results from the interaction terms are significantly different from zero, this suggests that the pre–treatment trends are not parallel, and that our estimates are most likely biased. The results are presented in Table 11 in Appendix A.3.3. For the three outcome variables on labor demand, we observe statistically significant results for both placebo treatments in 1995 and 1996, except for the the unemployment rate in 1996. These results are worrisome as they suggest that the trends deviate systematically in the pre–regulation period, which is a violation of the parallel trends assumption. To conclude, pre–trends might be driving our estimated results in Table 4.

Changes in the Working Age Population

Our DiD estimates of the unemployment and the employment rates are incompatible, as they movie in opposite directions. Therefore, we want to investigate whether there are changes in the working age population in the treated municipalities as a consequence of the environmental regulations. An example of this could be selective migration as a consequence of treatment. Investigating movements and treatment effects on the working age population will provide insight regarding the DiD results. The evolution in the working age population as a share of the total population for the treatment group and the control groups is depicted in Figure 12 in Appendix A.3.2, and seems to comply with the parallel trends assumption. In addition, the related DiD estimates with the share as the dependent variable are presented in Table 9. The DiD estimates prove a statistically significant negative effect on the working age population from the 2000–restriction. More specifically, the 2000–restriction is associated with a decrease of approximately 0.003 percentage point relative to the pre–treatment mean. This negative effect might indicate that our estimates of the employment rate and payroll tax are overestimated in Phase 2, implying that some factor affecting the working age population in the second phase drives the estimates on employment rate and payroll tax.

Selective migration could cause the decrease in the working age population after 2000 in the treated municipalities. The evolution in the emigration as a share of total population for the treatment and the control groups are depicted in Figure 13 in Appendix A.3.2, and also seem to comply with the parallel trends assumption. Nonetheless, the results in Table 10 show that the effects of emigration as a share of the total population are not statistically significant. Therefore, there must have been reasons other than emigration which have caused the changes in the working age population. Yet, it is beyond the scope of this study to investigate the reason to the decrease in the working age population

Omitting Treatment Municipalities Singly

With only seven municipalities in our treatment group, there is a risk that certain municipalities drive our results. The position of the aluminum plants as cornerstone companies varies across the treated municipalities. This, together with differences in the plants' technologies and the times of their Søderberg–shutdowns, lead us to suspect that the effect on labor demand may differ across the treated municipalities. Leaving out the treated municipalities singly from the treatment group enables us to observe whether the estimates change, and to evaluate the importance of each municipality in our results. The results of these robustness tests related to municipality dependency are presented in Tables 12–18 in Appendix A.3.4.

Overall, when excluding the municipalities singly from the DiD model, the results for the unemployment rate and employment rate are consistent with the previous findings of the full sample. This implies that the identification does not hinge on one municipality. The size of the point estimates and significance for Phase 2 interaction term differ slightly, but the differences are negligible in size. Additionally, both the preceding signs and the statistical significance of the results are overall similar.

The results for payroll tax per person are less consistent when excluding municipalities singly from the treatment group. The point estimates are relatively similar, although the significance level depends on which municipality we exclude. When excluding Høyanger, Sunndal, Farsund and Vefsn, the coefficient of $(Treat_m \cdot Phase1_t)$ is no longer significant. However, when excluding Karmøy or Kvinnherad, the estimates for Phase 1 are more significant. On the other hand, the results for Phase 2 are more consistent with previous findings. Yet, the results become more significant when excluding Farsund and Vefsn, while when excluding Årdal they are insignificant without control variables. Hence, when using the data on payroll tax, the identification seems to hinge more on each individual municipality in both phases, and we should interpret the results for this outcome variable with caution.

6.2 Main Results Instrumental Variables Approach

In the following section, we estimate the effect of a technological advancement proceeding of an environmental regulation on outcomes of labor demand. The analysis consists of four main parts. First, we present our first–stage estimates in order to evaluate the introduction of increased stringency in emission permits in 2000 as an instrument for the shutdown of the Søderberg anodes in the aluminum plants. Second, we run the instrumental variable regressions. Third, we present our reduced–form estimates, the indirect effect of the 2000–restriction on our outcomes of interest. Lastly, we run a sensitivity analysis to identify the weaknesses of our model.

6.2.1 First–Stage Estimates

Table 5 reports the estimates of the *first–stage* relationship between the period after the 2000–restrictions and the shutdown dates of the Søderberg anodes. More specifically, we estimate Equation 3 presented in Subsection 5.3, and exploit the 2000–restriction as a potential instrument to predict the shutdown of the Søderberg anodes. First, we regress *Søderberg* on *Post2000* without controls, followed by the regression including controls.

Table 5: First–Stage Estimates of Søderberg–shutdown on the 2000–restriction

	(1) Søderberg off Mean pre–treatment	(2) Søderberg off Baseline	(3) Søderberg off Baseline with Controls
	0		
Post2000		0.366***	0.363***
		(0.0546)	(0.0542)
Observations	1208	2869	2869
Adjusted R^2		0.362	0.364
Control Variables		No	Yes
Municipality Linear Trend		Yes	Yes

Notes: Each column represents results from separate OLS regressions based on Equation 3. The dependent variable is a dummy variable equal to one in the years following the shutdowns of the Søderberg anodes in the aluminum municipalities, and zero otherwise. The sample consists of 151 municipalities over the sample period 1992-2010. The control variables included are described in Section 5.3.1. Standard errors are in parentheses and are robust to heteroskedasticity.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

In Table 5, we observe that without controlling for other confounding factors, the 2000–restriction is associated with a 0.366 probability of a treated municipality experiencing a shutdown of Søderberg anodes, relative to a pre–treatment mean of zero. The instrument is robust to including controls as depicted in Column (3), as the probability remains almost identical. The first–stage estimates are consistent with what is indicated in Table 1 in Section 2: that all the shutdown dates are after the year 2000, and that an increase in the probability of a Søderberg–shutdown after 2000 is therefore anticipated. All first–stage estimations are statistically significant, as the test of whether the coefficient of *Post2000* is equal to zero yields an F–statistic above 770 when controls are not included and above 260 when all the controls are included, thereby higher than the requirement of 10.[28] Thus, the

[28]The criterion by Staiger and Stock (1997) is met.

instrument is relevant, and the first–stage estimates suggest that the years following the increased stringency of the emission permits in 2000 can be used to predict the shutdown of the Søderberg anodes in the treated municipalities.

6.2.2 Instrumental Variable Estimates

Columns (1) and (2) of Panel A in Table 6 present estimates from the OLS regression of the unemployment rate on Søderberg–shutdown. The OLS results are not consistent across specifications as they prove to be sensitive to the introduction of covariates. The 2SLS estimates reported in columns (3) and (4) of Table 6 point toward insignificant effects of the Søderberg–shutdown on the unemployment rate. Without including covariates, the coefficient points in the same direction as the OLS estimates. Moreover, the size of the effect is similar to the OLS estimate, but it is insignificant at the five percent level. When including the covariates, the estimates point in the opposite direction, and the Søderberg–shutdown is not statistically significant at any level. Hence, our estimates suggest that there is no significant effect of the Søderberg–shutdown on the unemployment rate.

The first two columns of Panel B in Table 6 present estimates from the OLS regression of the employment rate on the Søderberg–shutdown. As opposed to the 2SLS estimates of the unemployment rate, we observe that both the OLS and 2SLS estimates are less sensitive across specifications, and that they remain positive, as well as statistically significant at conventional levels. The size of the effect of the Søderberg–shutdown on the employment rate is greater when introducing the instrument in 2SLS estimation depicted in Column (3), and is still statistically significant. However, when introducing the covariates, the size of the effect decreases from 0.12 to 0.04 percentage points. Specifically, this represents an increase of approximately eight percent compared to the pre–2000 mean. As the estimates are fairly consistent across specifications, and the preceding sign remains, we can conclude on a positive effect of the Søderberg–shutdown on the employment rate in the treated municipalities.

Table 6: OLS and 2SLS Estimates for Outcomes of Labor Demand

	(1) OLS	(2) OLS	(3) 2SLS	(4) 2SLS
A. Dependent variable: Unemployment Rate				
Søderberg off	-0.00736***	-0.000356	-0.00817*	0.00255
	(0.00230)	(0.00220)	(0.00421)	(0.00394)
Pre-2000 mean	0.0363	0.0363	0.0363	0.0363
Observations	2718	2718	2718	2718
F-test of instrument			771.11	260.21
B. Dependent variable: Employment Rate				
Søderberg off	0.0728***	0.0236***	0.120***	0.0447***
	(0.00406)	(0.00672)	(0.0174)	(0.0145)
Pre-2000 mean	0.538	0.538	0.538	0.538
Observations	2869	2869	2869	2869
F-test of instrument			814.25	274.47
C. Dependent variable: Payroll Tax				
Søderberg off	-58.17	-3541.3***	-2668.9*	-5299.7***
	(555.7)	(751.6)	(1376.6)	(1152.2)
Pre–2000 mean	12709	12709	12709	12709
Observations	2805	2805	2805	2805
F–test of instrument			795.86	268.33
Control Variables	No	Yes	No	Yes
Municipality Linear Trend	Yes	Yes	Yes	Yes

Notes: Each column represents results from separate OLS regressions based on Equation 4. The sample consists of 151 municipalities over the sample period 1992-2010. The dependent variables are the unemployment rate in Panel A, the employment rate in Panel B and the payroll tax paid by each person in the working age population in Panel C. The control variables included are described in Section 5.3.1. The pre–2000 mean is the mean of the dependent variable before 2000. Standard errors are in parentheses and robust to heteroskedasticity.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The results from the OLS regression of the payroll tax per person in the working age population on the Søderberg–shutdown are depicted in Columns (1) and (2) in Panel C. The OLS estimates differ largely in size, and in statistical significance, across specifications. The introduction of the instrument is depicted in the 2SLS estimates in Columns (3) and (4). Without including control variables, the effect from the Søderberg–shutdown is insignificant at the five percent level. When including controls, however, the estimates indicate a larger decrease in payroll tax of 5,299.7 NOK associated with the Søderberg–shutdown. This estimate is statistically significant at conventional levels. Conclusively, the results do not prove consistent across specifications. The 2SLS estimate in Column (3) appear upward biased, meaning it is associated with the labor market characteristics women in the working force, educational attainment, and the aluminum price. Thus, the 2SLS estimate in column (3) remains endogenous.Because as OVB plays a large role in the estimate in column (3), it is likely that there exist other crucial unobservable confounding factors we are not in position to control for. We cannot know if the estimate in Column (4) is upwardly or downwardly biased. Hence, we conclude that the estimated effect of Søderberg–shutdown on the payroll tax per person in the working age population must be interpreted with high vigilance. Our findings will further be discussed in Section 7.

6.2.3 Reduced–form Estimates

Table 7 presents the reduced–form estimates, based on Equation 5, which is the direct regression of the three outcomes of interest on the 2000–restriction. The advantage of the reduced–form approach is that it does not rely on excluding the 2000–restriction indicator from the second–stage equation. Panel A suggests that the causal relationship between the 2000–restriction and the unemployment rate is weak or nonexistent (Angrist & Pischke, 2015, p. 146), which supports our conclusion in Subsection 6.2.2. Panel C shows that the estimated effect on the payroll tax vary to a large extent across specifications. As a result, it points toward more OVB not accounted for in the model, as in the second–stage estimates. Therefore, drawing inference regarding the unemployment rate and the payroll tax per person must be handled with strong caution. On the other hand, the statistically significant reduced–form estimate of the employment rate of 0.016 percentage points, and the consistency across specifications, support the findings that the Søderberg–shutdown increased the employment rate. To conclude, the only causal effect of interest of the IV results is the one estimated for the employment rate.

Table 7: Reduced–form Estimates for Outcomes of Labor Demand

	(1)	(2)
A. Dependent variable: Unemployment Rate		
Post 2000	-0.00299*	0.000927
	(0.00158)	(0.00142)
Pre–2000 mean	0.0363	0.0363
Observations	2718	2718
Adjusted R^2	0.064	0.212
B. Dependent variable: Employment Rate		
Post 2000	0.0440***	0.0162***
	(0.00551)	(0.00499)
Pre–2000 mean	0.538	0.538
Observations	2869	2869
Adjusted R^2	0.063	0.616
C. Dependent variable: Payroll Tax		
Post 2000	-977.2**	-1923.5***
	(458.3)	(414.6)
Pre–2000 mean	12709	12709
Observations	2805	2805
Adjusted R^2	0.126	0.359
Control Variables	No	Yes
Municipality Linear Trend	Yes	Yes

Notes: Each column represents results from separate OLS regressions based on Equation 5. The dependent variables are the unemployment rate in Panel A, the employment rate in Panel B and the payroll tax by each person in the working age population in Panel C. The sample consists of 151 municipalities over the sample period 1992-2010. The covariates included are described in Section 5.3.1. Standard errors are in parentheses and robust to heteroskedasticity.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

6.2.4 Sensitivity Analysis of the IV estimation

Test of Endogeneity

In order to confirm that the *Søderberg* variable in fact is endogenous, we run the Durbin–Wu–Hausman test of endogeneity. The results of the test are depicted in Appendix A.3.5. With a p–value at conventional levels, we reject the null hypothesis that the residual is insignificant in the OLS regressions of the unemployment rate, the employment rate, and the payroll tax on *Søderberg*. Consequently, we confirm the relevance assumption entailing that *Søderberg* is an endogenous variable. The results are consistent across specifications, except for a minor increase in the p–value for the unemployment rate when including control variables.

Extended Sample

As mentioned in Section 4, we extend the sample in numbers of years from the last year of sample from 2010 to 2013, to explore the potential to capture the effects from the shutdown of the Sødeberg anodes in Karmøy in 2009. Because our main sample has observations until 2010, we fear that we are missing the effect from the last Søderberg–shutdown. The 2SLS estimates are presented in Table 21, first–stage estimates in Table 20 and reduced–form estimates in Table 22 in Appendix A.3.6.

The estimates and significance in the first–stage of the extended sample are similar to our main sample. Table 20 in Appendix A.3.6 reports significant coefficients of the Søderberg anode shutdown, which indicates that the instrument still meets the relevance condition.

Column (3) and (4) in Table 21 in Appendix A.3.6 reports the 2SLS estimates with the extended sample. Overall, the second-stage estimates change little when expanding the sample period to 2013. All of the outputs are reported with the same preceding sign as in Table 6. The significance levels vary only slightly, and so does the size of the point estimates. The estimates of the unemployment rate and of the the payroll tax, presented in Panel A and in Panel C respectively, are still not robust across specifications. However, the estimates of the employment rate presented in Panel B, are still significant and robust across specifications. Only the size of the point estimates with covariates are reduced to half compared with our previous findings. Thus, in examining the estimates of the employment rate, expanding the sample by three additional years, decreases, rather than enlarges, the effect from the Søderberg–shutdown on the employment rate.

Lastly, the reduced–form estimates with the expanded sample presented in Table 22, are similar to the reduced–form results presented with the sample up to 2010. The only difference is that the significance level is reduced in the estimated effect of the Søderberg–shutdown on the employment rate when including the control variables and the municipality–specific linear trend. To summarize, the conclusions are similar to the findings presented in Section 6.2.2, except that the size of the effect on the employment rate diminishes in the second–stage with an extended sample. Hence, the effects from the Søderberg–shutdowns seem to be short–term.

6.3 Summary of the Results

The results from our DiD model suggest that the short–term effects (i.e., Phase 1), of the 97–agreement on the unemployment rate and the payroll tax per person in the working age population, are insignificant at a five percent level. Our findings suggest a significant 2.6% increase in the employment rate in Phase 1. The estimates of the unemployment rate, the employment rate and the payroll tax per person pass the checks for anticipation effects. The first two also pass the check for municipality dependency in the estimates in Phase 1. However, the estimates of the payroll tax hinge on the inclusion of certain municipalities, implying that some municipalities are more important for the estimates of this outcome variable than others.

Further, our DiD estimations show a significant decrease in the unemployment rate, the employment rate, and the payroll tax paid per person in the working age population, in the years subsequent to the 2000–restriction (i.e., Phase 2). More specifically, the 2000–restriction led to a 15.5%, 1.7%, and 11.8% decrease, respectively. That the estimates of the employment rate and the payroll tax have matching preceding signs, indicates that the income of the treated municipalities decreased after the 2000–restriction and that less people worked. A sensitivity check shows a significant decrease in the working age population after 2000, which have caused an overestimation of employment rate and payroll tax per person, in Phase 2. Yet, the results are still reliable. The unemployment rate estimates pass all the sensitivity checks in Phase 2.

The identifying assumptions of the DiD model seem to be met (i.e., the parallel trends assumption and the no–anticipation assumption) when examining the trend plots in Appendix A.1 and checking for anticipation effects. The placebo test in Subsection 6.1.1, however, points to the existence of underlying trends in our outcomes of interest. Consequently, the parallel assumption could be violated.

The results from our IV model in estimating the effect from a technological advancement, the Søderberg–shutdown, indicate a 8.3% increase in the employment rate, which is statistically significant. The effect of Søderberg–shutdown on unemployment rate is insignificant. The estimates of the payroll tax per person in the working age population, differ substantially across specifications in second–stage and reduced–form. As a result, this implies that that there exist other crucial unobservable confounding factors we are not in position to control for, and we cannot know if the estimates will be upward or downward biased. Hence, the only significant effects found are the ones for the employment rate. The Durbin–Wu–Hausman test of endogeneity confirms that our instrumented variable is endogenous, and the estimates for the employment rate continue to be consistent when we extend the sample period.

7 Discussion

In the following section, we discuss our estimated results and potential drawbacks. Further, we bring up limitations to the dataset, limitations imposed on our study by virtue of our estimation strategies, as well as examining the external validity of our study. To conclude, we discuss how our study contributes to existing literature, as well as recommendations for further research.

7.1 Discussion of the Results

Our findings suggest that environmental regulations lead to short–term increases in employment. As described in Section 3, with a model on labor demand under environmental regulations, Berman et al. (2001) declare that the true dominating effect of environmental regulation cannot be predicted from the model alone. Whether the effect is positive, or negative, will depend on the type of abatement activity introduced as a result of environmental regulations, and whether or not this activity is a substitute, or a complement, to labor (Berman & Bui, 2001; D. Autor, Levy, & Murnane, 2003). According to the reviewed literature (D. Autor et al., 2003), the true consequences of a technological change originating of environmental regulation should be evaluated in relation to the tasks replaced by the pollution abatement technology.

The estimates of our DiD model show that the effect on the unemployment rate of the 97–agreement is insignificant in Phase 1. However, we cannot observe if there might be short periods with frictional unemployment, or if workers for instance experience a transition from part–time employment to full–time employment. In Phase 2, there is a significant decrease in the unemployment rate. Hence, there are fewer people categorized as unemployed. Abatement activities may have complemented workers (D. Autor et al., 2003) through new patterns of activities in line with the "induced" innovation hypothesis (Jaffe et al., 2002; Ashford & Heaton, 1983).

The DiD results suggest a positive effect of the environmental regulation on the employment rate in the first three years after the 97–agreement. The increase found, may be a result of the abatement activities introduced to comply with the regulation, complemented labor, similar to what we see in Phase 2 for the unemployment rate. The 2000–restriction entailed that the halls with Søderberg anodes had to be shut down by 2007. Thus, we would expect a decrease in the demand for labor resources because the people employed in the halls would lose their jobs, to prepare for the Søderberg–shutdown. Our DiD estimates in Phase 2 on the employment rate comply with these expectations. However, in the IV estimates, there is a positive effect of the technological advancement on the employment rate. Consequently, the released labor resources from the aluminum industry seem to have been allocated to other industries, in line with what is proposed for competitive markets by Mabry et al. (1986).

The main reason to expected changes in labor demand in Phase 2 in DiD is the Søderberg–shutdown. Our IV model pinpoints this, by instrumenting the 2000–restriction on the Søderberg–shutdown. The positive estimated effects of the Søderberg–shutdown in the IV outputs, are larger than the estimated effects of the 2000–restriction in the DiD outputs. We will be careful to conclude on a comparison of these two models, as both bear limitations. However, one possibility is that by pinpointing our estimations of interest to the Søderberg–shutdown in IV, the negative effect found in Phase 2 for the employment rate might be dominated by the long–term effects of the 97–agreement. Thus, the effect found in Phase 2 is not necessarily only caused by the 2000–restriction, and we get back to this in Subsection 7.3. Moreover, when expanding the IV sample with three more years, we find that the effect of Søderberg–shutdown on the employment rate diminishes rather than enlarges. This indicates that the effects of the Søderberg–shutdown are short–term. Thus, we find that short–term effects of the 97–agreement and the Søderberg–shutdown, as a result of the 2000–restriction, are positive on the employment rate.

In addition, we have to keep in mind that our results on the employment rate might be overestimated after 2000, as the denominator, the working age population, decreases in this phase. This decrease may overestimate both the Phase 2 results in DiD and the IV results. The size of the effect on the employment rate from Phase 2 in DiD is five times larger than the effect of the Søderberg–shutdown. Therefore, despite a potential overestimation, the positive effect of Søderberg–shutdown on the employment rate is still likely. Thus, we are confident to conclude on the short–term effects, but not the long–term effects of each of these environmental regulations on the employment rate.

The payroll tax per person in the working age population is a proxy for the generated income in the municipality. It is expected to mirror the effects in the employment rate, which it does in our estimated results. Yet, the payroll tax is differentiated across zones, which implies that the aluminum municipalities (the treatment group) are subject to different rates of payroll tax. This accords with the fact that the estimates of payroll tax per person are sensitive to omitting municipalities singly, which gives a hint of heterogeneous effects across the municipalities. Thus, we fear the presence of OVB in our DiD estimations for the payroll tax which we are not able to control for. Also, in these estimates, we have to take into consideration the decrease in working age population. One last point of consideration, as mentioned in Section 6.1.1, there is a chance that there are underlying trends driving the DiD results on the three outcome variables of interest.

The DiD estimates of the employment rate and the unemployment rate move in the same direction in the second phase, the years after 2000, and this is a cause for inquiry. We would have expected these rates to move in opposite directions as a ratio to the labor force. Nonetheless, the four margins mentioned in Section 4.1 indicate that these two rates are not entirely symmetric, as the environmental regulations may have affected labor demand dissimilar through these margins. Moreover, there are mechanisms related to hidden unemployment not absorbed by the unemployment rate. It is beyond the scope of this study to distinguish between the different mechanisms because we only observe the net movements in the unemployment rate and the employment rate on the municipality–level.

7.2 Limitations to the Dataset

7.2.1 Municipality–level Data versus Plant–level Data

The changes in the unemployment rate, the employment rate, and the payroll tax per person at municipality–level may, or may not, depict changes in plant–level employment. Access to plant–level data like Walker (2011) and the direct abatement cost related to compliance for each plant, would provide more details on the gross movements in the aluminum plants from the environmental regulations, as a complement to our estimated net effects. However, the fact that the aluminum plants are cornerstone companies, and essential employers in the associated municipality (Godal, 1998), arguably validates our data at the municipality–level as a measure of the net effects in the municipalities.

7.2.2 The Working Age Population as Denominator

To use the working age population as the denominator in the employment rate, as well as in the payroll tax paid per person in the working age population, can be misleading, and is a potential drawback. In the sensitivity analysis, we find that the change in estimates of the working age population is negative and significantly different from zero after 2000. Evolution in the size and composition of the working age population may drive the outcomes of interest up or down. The working age population consists of both people inside and outside the labor force, indicating that a decrease in the employment rate and payroll tax can be a result of multiple scenarios. First, a worker may transition from being employed to unemployed, or to outside the labor force. Moving outside the labor force would be the case if the person, for instance, becomes incapable of working, retires, or goes back to study. Second, the total working age population may increase because more people move into the category, either because they turn 16 years old, or because of more immigration into one of the respective municipalities. We do not have access to the accurate measurement of the labor force, according to LFS (SSB, 2019a). Thus, utilizing the working age population as a denominator in the employment rate and the payroll tax per person enables us to capture some of the mechanisms of people reallocating in the working age population, both inside the labor force or outside the labor force. Still, a weakness of this variable chosen to enable comparison is that these ratios carry the risk of being biased by changes in the denominator (i.e., the changes in the working age population found in Subsection 6.1.1).

The unemployment rate is difficult to measure, as persons receiving disability benefits, sickness absence, or some other kind of public transfer income, are not defined as unemployed. Thus, these persons amount to the pile of hidden unemployment. Moreover, there may be several reasons for a worker to end up unemployed, either the person loses her job, leaves her job, is trying to re–enter the market after spending time outside the labor force, or is a new entrant as a result of for example recently graduating from high school or college (i.e., frictional unemployment) (Borjas, 2016). Also, being unemployed over time may cause persons who would like to work to withdraw from the labor force because they cannot find jobs they are qualified for (i.e., structural unemployment). Therefore, it could cause discouraged workers not to be included in the unemployment rate. This demonstrates that the group of unemployed persons is diverse, and makes it hard to say exactly what the change of the 2000–restriction on unemployment involves.

Given the discussion above, the employment rate and the unemployment rate are difficult to compare, as their denominators differ. More precisely, the labor force is defined as the denominator in the unemployment rate of LFS.[29] In the employment rate and payroll tax per person, the denominator is the working age population, which lumps together the whole population between age 16 and 74 (i.e., the unemployed, employed and those outside the labor force). Therefore, the change in the employment rate is less likely to be symmetric to the change in unemployment. This points to the importance of movements across different margins. We would need individual data to say exactly what the re–allocations in the labor market look like. However, even with this lower level of data, the categorization and grey zones of unemployment (i.e., the hidden unemployment) will make it hard to capture the true effects empirically.

[29]Labor Force Survey conducted by SSB. See Section 4.1 for more information.

7.2.3 Controlling for Differences in Industry Composition

Ideally, we would have liked to control for differences in the industry compositions in municipalities, because structural differences in unemployment across sectors might drive our results. A limitation in our data is the lack of a variable representing the share of businesses in the service sector.[30] We are not able to investigate this variable, nor possibly control for it in case of differences between control and treatment group.[31] Thus, there may be an issue with OVB related to this characteristic in our estimated results.

7.2.4 Controlling for Changes in Electricity Price

The fact that the aluminum industry is sensitive to the electricity price makes the price of the long–term contracts a potential factor that could drive the changes in labor demand. The electricity spot price, which is the market price of electricity, is likely to be effected by CO_2–quota prices set by the EU ETS system, and changes would gradually happen in the long–term contracted electricity price (Jæger & Budalen, 2019). The aluminum plants experienced an increase in the electricity price to the double from 2000 to 2011 (Magnussen, Spilde, & Killingland, 2011). The electricity price faced by the plants is time–invariant in the short–term (i.e., because of the long–term contracts) and municipality specific (i.e., each plant has its own contract with the local producer of electricity). Thus, a combination of the municipality fixed effects and year specific effects should have captured this increase in price.

7.3 Limitations to the DiD Estimation Strategy

7.3.1 Underlying Time Trends

One of the main premises of the DiD strategy is that the parallel trends assumption must hold, as discussed in Section 5.1.3. Based on the figures of the development of our outcome variables of interest in Appendix A.1 and A.3.1, we find that the trends before 1997 are visually and statistically similar between the treatment and control group. However, we can not guarantee that the aluminum plants never would have conducted abatement activities voluntarily in the absence of the 97–agreement and the subsequent

[30] As mentioned in Subsection 4.2 there is typically lower unemployment in areas with a high share of businesses in the service sector.

[31] We encountered the problem to be that big enterprises in Norway often have their main offices in Oslo, and in alternative databases on venture level, it is not possible to obtain data on for instance Norsk Hydro's plants in Sunndal and Årdal, because the company is registered in Oslo (SSB, 2014).

2000–restrictions, which would cause a violation of the parallel trends assumption. Due to increased overall focus on the improvement of environmental quality, we suspect that eventually, the aluminum industry would have to give way for national, and international, pressure regarding this awareness. As mentioned in Section 2, the CO_2–tax was already introduced in 1991, and with debate on expanding this, or a similar reform, to include all Norwegian industries including the aluminum industry, gives us reason to believe that the market players in our industry of interest anticipated the introduction of some policy. Moreover, their active participation in lobbying also supports this reasoning. In our analysis, we do argue that the timing of the conducted abatement activities would not have been the same as after the 97–regulation and the 2000–restrictions. However, in Subsection 6.1.1, it is suggested that the trends of our outcome variables of interest are systematically deviating in the pre–regulation period, which could give a hint that the discussed scenario may have occurred.

7.3.2 No Clear Cutoffs

To facilitate the application of a DiD method, there should be a sharp change in economic environment or changes in government policy for this method to provide evidence of a causal relationship (Arrow & Intriligator, 1999, p. 1296). As mentioned in Section 2, Vevatne et al. (2004) indicate that a voluntary agreement may have had little effect on the ability to compete and pollution level, which points to that 1996/1997 is not a clear cutoff. The long compliance period debilitates 1996/1997 as cutoff, yet, by dividing our period in two, we correct for this issue. Because the agreement in 1997 was voluntary, the penalties may have been too weak for the change and abatement activities to be implemented fast. Thus, it is less surprising that we do not find significant effects from the 97–agreement in Phase 1.

The 2000–restriction seems to have taken the aluminum industry more by surprise. Thus, 1999/2000 is arguably a more exogenous cutoff, indicating that the aluminum industry would not have implemented means to reduce their emissions in the same from 2000 to 2010. Yet, if there is a long–term effect of the 97–agreement happening after 2000, it is impossible to segregate such an effect of the 2000–restriction. Hence, we cannot draw inference on which of the environmental regulations that cause the effects that we observe in Phase 2.

7.3.3 Small Treatment Group

Despite the total number of observations in our sample being relatively large, there is a small number of treated municipalities. Traditional inference assumes that the number of observations is large. Thus, the small size of our treatment group may cause imprecise estimate, and less accurate test statistics. In the figures of validation of the differential evolution in Appendix A.3.1, especially in Figure 9, that shows the differential evolution in the unemployment rate, we observe relatively large confidence intervals, which may be caused by the small treatment group. The fact that we find significant results despite the number of treatment municipalities being small, might be a reason to doubt that we have evidence to draw an inference and to believe that our estimates are either biased or that we are estimating the effect of an underlying trend. Another possible consequence of the small treatment group is that our findings cannot be generalized and used as an outlook on what would have happened to other sectors or businesses. Yet, we argue they can, as the seven municipalities making up our treatment group represent the entire primary aluminum industry in Norway.

7.3.4 Other Mechanisms Driving the Results

In addition to the mechanisms mentioned in Section 7.2, other mechanisms may drive changes in all our variables of interest, and lead to biased results. As described in Subsection 6.1.1, selective migration could be a problem in our setting, as decreased employment coming from abatement activity could lead to people migrating, which again could bias our results upwards. Moreover, the 97–agreement may have led to more high school graduates deciding to attain higher education in the years posterior to 1997, because they anticipate that the local aluminum industry will need less workers in the future. If this happened in reality, we would have a problem of collinearity and bias in the model. In spite of that possibility, there seems to be no increase in the educational attainment of the treatment group relative to the control group looking at Table 3. On the contrary, the only evolution in the differences in educational attainment we find between the two groups, is a slight increase in the control group's share of adults attaining higher education up to four years. However, this is negligible. Another possible mechanism could be that if plants can abate at low cost, the employment effects are smaller (Berman & Bui, 2001). If this is the case, our estimates of the effect of environmental regulation on outcomes on labor demand will could be underestimated.

7.4 Limitations to the IV Estimation Strategy

7.4.1 Few Municipalities Experiencing Søderberg–shutdown

The treatment group experiencing treatment consists of few observations, that is that only five of the seven aluminum plants experience a shutdown of Søderberg anodes. This may cause difficulties in the interpretation, and isolation, of the effect of the technological advancement on labor demand outcomes, due to the small variation in our independent variable of interest. Hence, the estimated effect of the Søderberg–shutdown on the employment rate, cannot explain the 2000–restrictions' effect on the employment rate in the two municipalities that did not experience the shutdown, or whether there is one. This effect will only be captured in our DiD model. If there is one, it would be a violation of the exclusion restriction, because *Post2000* would have an effect through other channels than the Søderberg–shutdown on the outcomes of interest.

7.4.2 Endogenous Instrument

Ultimately, we fear that the 2000–restriction is not entirely exogenous as an instrument, because it may be correlated with the error term in our IV model. If the instrument is endogenous, the independence assumption is violated, and our results will be biased. The 2000–restriction only affected the municipalities with aluminum industries, that is to say, municipalities that are likely to have a labor market, affected by the primary aluminum plant activity. Thus, they are fundamentally different from the other municipalities in our sample. In our DiD estimation, we have controlled for differences in municipalities by comparing the treatment group with a comparison group, as they can never be argued to be completely equal. Yet, no matter how many control variables we include in the IV estimation, we will have to assume that the instrument is entirely unrelated to unobservable OVB that might affect labor demand. We can never be fully able to control for the comparison of fundamentally different municipalities in our IV analysis. Consequently, the independence assumption is arguably violated.[32]

[32] To solve this issue, we could decrease our sample to only consisting of the treatment municipalities. Nonetheless, this would make inference highly unreliable, as the sample would consist of few observations.

7.5 Implications of Our Study

This study contributes to the existing literature on the effect of environmental regulations on labor demand (i.e., Henderson, 1996; Greenstone, 2002; Berman & Bui, 2001; Ferris et al., 2014; Liu et al., 2017; Walker, 2011; Walker, 2013), and about the effect of technological change on these outcomes (i.e., Mabry & Sharplin, 1986; D. Autor et al., 2003; D. H. Autor, 2013). The findings are arguably relevant in current national and international policy debates regarding environmental protection and are beneficial for policy–makers when evaluating the implementation of environmental regulations.

Central to any study is the extent to which the results are possible to generalize to other situations and contexts, more specifically, whether they are externally valid. We choose to estimate the net effect of environmental regulations to enable the application of municipality–level labor market outcomes as proxies for the labor demand in the aluminum industry, as opposed to for instance Greenstone (2002), who calculates the gross effect. The estimated effects can be directly transferable to other countries with industries meeting similar stringency in emission permits or other types of new penalties. However, it will depend on the characteristics of the industry, how labor intensive it is, its impact on the local community, and other municipality and plant characteristics. We argue that our estimated effects are transferable to municipalities, or areas, depending on one main employer, located in remote areas. Moreover, the effects can be transferred to the secondary aluminum industry and the overall GDP of the Norwegian manufacturing industry. Based on this discussion, it would be interesting to analyze whether similar results could be obtained studying more industries in Norway, or in other countries. The estimated employment effects should be of interest to regulators elsewhere, or in other industries, due to the current national and international policy debates regarding environmental protection.

7.6 Further Research

This thesis is limited to studying the unemployment rate, employment rate and payroll tax per person in the working age population, at the municipality–level. For other studies, it would be interesting to investigate additional outcome variables, heterogeneous effects in the outcome variables, as well as to use individual–level data or plant–level data. In particular, it would be of interest to investigate whether unemployment and employment differ across sectors, or across demographics. This could enable studies on the costs of labor reallocation related to regulation, which Walker (2011; 2013) claims that are more important than the effects on the level of employment.

With available data on plant–level one could explore the plant–level change in employed persons, destruction of jobs, and the costs faced related to abatement activities or penalties. Besides, as proposed by Popp et al. (Popp et al., 2010), another possible result which is interesting to examine, is the effect on input factors other than labor (i.e., capital and natural resources). Further, it could be interesting to investigate potential spillover effects to other actors, both commercial and non–commercial. Thus, it would be of interest to investigate effects on other actors in the local community of the municipality, in neighboring municipalities or in other parts of the supply chain situated other places in the country or abroad. More specifically, the environmental regulations could affect a local non–profit organization (e.g., the local sports clubs common in Norway, and charity), the secondary production of aluminum located in other parts of the country, or new establishments, like proposed has occurred in Årdal (Johansen, 2008).[33] Ultimately, with the caveat of significant pre–trends in our estimations in mind, it would be of interest to investigate the drivers behind these trends in future research.

[33]Norwegian culture for volunteer engagement hinges on contribution from businesses.

8 Conclusion

This thesis aims to analyze the research question "*How did the environmental regulations, namely the 97–agreement and the 2000–restriction, and the related technological advancement, affect local labor demand in the municipalities which encompass a primary aluminum plant?*". To answer the research question, we have used a DiD approach with two phases. With our DiD strategy, we examine the causal effect of the 97–agreement, the voluntary agreement signed between the primary aluminum industry and the Ministry of Climate and Environment in 1997, and the 2000–restriction, the increased stringency of the emission permits of the primary aluminum industry in 2000, on labor demand. Further, we use an IV approach to investigate the effect of the technological advancement, the shutdown of the more polluting, and less efficient, Søderberg anodes between 2003 and 2009, on labor demand. To do this, we utilize data collected from the Norwegian Centre for Research Data and Statistics Norway that include municipality and labor characteristics on a municipality–level.

Our robust and precise DiD estimates show that the 97–agreement increased the employment rate in the short–run in the associated municipalities, the municipalities in which the primary aluminum plants are located. In aggregated terms, the 97–agreement resulted in approximately 135 extra jobs the first three years in each of the seven municipalities, relative to the pre–treatment mean. The long–term effects of the 97–agreement and 2000–restriction are hard to separate because they overlap. Yet, the main purpose of the 2000–restriction was to enhance a shutdown of the Søderberg anodes. Using our IV model, instrumenting the 2000–restriction on the Søderberg–shutdown, the results show a positive short–term effect of the Søderberg-shutdown on the employment rate. In aggregated terms, the Søderberg–shutdown resulted in approximately 406 extra jobs in each of the five associated municipalities, relative to the pre–2000 mean. Thus, our DiD and IV results together show positive short–term effects of the environmental regulations and the technological advancement associated with the environmental regulation, on the employment rate. Our DiD estimates on the employment rate also show that the phase from 2000 to 2010 is related to a destruction of 90 jobs in the seven associated municipalities, relative to the pre–treatment mean. This implies that the combination of the 2000–restriction and the long–term 97–agreement effects, on the employment rate, are negative. One of the mechanisms driving the results in the latter phase of our DiD sample, from 2000 to 2010, is the decrease in the working age population. The DiD estimates on our proxy for the size of the generated income, mirror, and thereby support, the results found for the effect on the employment rate in the latter phase of the DiD model. Our DiD results show no short–term effect of the 97–agreement on the unemployment rate. However, the environmental regulations, the combination of the 2000-restriction and the long–term effects of 97–agreement, led to a significant decrease in the unemployment rate after 2000 of 15.5%. From the placebo test, there is a caveat of potential underlying trends. In addition, we recognize the possibility that our instrumental variable is endogenous.

Overall, our results suggest that the short–term effect of the 97–agreement and the technological advancement associated with the 2000–restriction, is a significant increase in the employment rate. Moreover, the combination of the effect of the 2000–restriction, and the long–term effect of the 97–agreement (i.e., after the year 2000), result in a negative effect both on the employment rate and the unemployment rate. Hence, this study contradicts what the Norwegian media indicated about environmental regulation raising unemployment and threatening local communities. In addition to health benefits and higher quality of life for citizens through reductions in GHG–emissions, environmental regulations induce increased allocation of labor resources in the short–term. This study adds to the economic literature on the causal effect of environmental regulation on labor demand, and of technological advancement as a result of environmental regulation on labor demand, and represents the first such study in the context of Norway or other Nordic countries. Future research should focus on studying the consequences of further environmental regulations and associated technological advancement, on the drivers behind the underlying trends in outcomes of labor demand, and on identifying the effects of environmental regulations on individual–level outcomes.

References

Abadie, A., Athey, S., Imbens, G. W., & Wooldrigde, J. (2017). When should you adjust standard errors for clustering? *NBER Working Paper Series*, *C21*(24003), 1–18.

Acemoglu, D. & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics* (Vol. 4, pp. 1043–1171). Elsevier.

Alcoa. (n.d.). Norway. Retrieved May 19, 2019, from https://www.alcoa.com/norway/no/default.asp

Alfsen, K. H. (1998). Kyoto-avtalen: Hva kan norge gjøre? *CICERO Report*, (1).

Alfsen, K. H. (1999). Flexible instruments in climate policy. *CICERO Policy Note*, (1).

aluminum-production.com. (2019). Process basics. Retrieved May 14, 2019, from http://www.aluminum-production.com/process_basics.html

Andresen, S. & Butenschøn, S. H. (2001). Norwegian climate policy: From pusher to laggard? *International Environmental Agreements: Politics, Law and Economics*, *1*(3), 337–356.

Angrist, J. D. & Krueger, A. B. (1999). *Empirical strategies in labor economics*. Elsevier.

Angrist, J. D. & Pischke, J.-S. (2009). *Mostly harmless econometrics*. Princeton.

Angrist, J. D. & Pischke, J.-S. (2015). *Mastering 'metrics*. Princeton.

Arrow, K. J., Cropper, M. L., Eads, G. C., Hahn, R. W., Lave, L. B., Noll, R. G., ... Smith, V. K., et al. (1996). Is there a role for benefit-cost analysis in environmental, health, and safety regulation? *Science*, *272*(5259), 221–222.

Arrow, K. J. & Intriligator, M. D. (1999). *Handbook of labor economics*. Elsevier.

Ashford, N. A. & Heaton, G. R. (1983). Regulation and technological innovation in the chemical industry. *Law Contemporary Problems*, *46*(3), 109–57.

Autor, D. H. (2013). The "task approach" to labor markets: An overview. *Journal for Labour Market Research*, *46*(3), 185–199.

Autor, D., Levy, F., & Murnane, R. J. (2003). The skill content of recent technology change. *Quarterly Journal of Economics*, *118*(4), 1279–1333.

Barber, M. & Taberaux, A. T. (2014). The end of an era for søderberg technology in north and south america. *Light Metals*, *68*.

Beerli, A. & Peri, G. (2015). The labor market effects of opening the border: New evidence from switzerland. *NBER Working Paper Series*, (21319).

Berge, J. (2008, November 26). Jens har ikke gjort det han lovet. Retrieved May 14, 2019, from https://www.nettavisen.no/okonomi/--jens-har-ikke-gjort-det-han-lovet/2428820.html

Berman, E. & Bui, L. T. M. (2001). Environmental regulation and labor demand: Evidence from the south coast air basin. *Journal of Public Economics*, *79*, 265–295.

Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economic s*, 249–275.

Bø, T. P. & Næsheim, H. (2015). Hvorfor ulike arbeidsledighetstall? Retrieved April 11, 2019, from https://www.ssb.no/arbeid-og-lonn/artikler-og-publikasjoner/hvorfor-ulike-arbeidsledighetstall--236723. (Statistics Norway)

Borjas, G. J. (2016). *Labor economics*. McGraw-Hill Education.

Bruvoll, A. & Bye, B. (2009). Evaluering av effekter av virkemidler i klimapolitikken. *Statistics Norway*, (48).

Bye, T., Larsson, J., & Døhl, Ø. (1999). Klimagasskvoter i kraftintensive næringer: Konsekvenser for utslipp av klimagasser, produksjon og sysselsetting. *Statistics Norway*, (24).

Chay, K. & Greenstone, M. (2003). The impact of air pollution on infant mortality: Evidence from geographic variation in pollution shocks induced by a recession. *Quarterly Journal of Economics*, *118*, 1121–1167.

Cover, M. (2011, June 10). Economic study shows epa regulations increase prices, kills jobs. Retrieved May 14, 2019, from https://www.cnsnews.com/news/article/economic-study-shows-epa-regulations-increase-prices-kill-jobs

Dupuy, D. (1997). Technological change and environmental policy: The diffusion of environmental technology. *Growth and Change*, *28*, 49–66.

Ferris, A. E., Shadbegian, R. J., & Wolverton, A. (2014). The effect of environmental regulation on power sector employment: Phase 1 of the title iv so2*tradingprogram*. *Journal of the Association of Environmental and Resource Economists*, *1*(4), 521–552.

Førde, T. (2011, May 19). Norges største riveprosjekt. Retrieved May 19, 2019, from https://www.aftenbladet.no/okonomi/i/grV0L/Norges-storste-riveprosjekt

Gärtner, M. (2016). *Macroeconomics*. Pearson.

Godal, O. (1998). Metallindustrien i norge: Økonomi, sysselsetting og utslipp av klimagasser. *CICERO Policy Note*, (2).

Government.no. (n.d.). Moce. Retrieved May 13, 2019, from https://www.regjeringen.no/en/dep/kld/id668/. (Ministry of Climate and the Environment)

Gram, T. (2008, December 11). Fakta: Søderberg historie etter 90 år. Retrieved April 3, 2019, from https://www.tu.no/artikler/fakta-soderberg-historie-etter-90-ar/259496

Greenstone, M. (2002). The impacts of environmental regulations on industrial activity: Evidence from the 1970 and 1977 clean air act amendments and the census of manufacturers. *Journal of Political Economy*, *110*(6), 1175–1219.

Groven, K., Lundli, H.-E., & Aall, C. (1999). *Local klimapolitikk* (Prosjektrapport No. 4/99). Vestlandsforskning. Sogndal, Norway.

Havnes, T. & Mogstad, M. (2011). No child left behind: Subsidized child care and children's long-run outcomes. *American Economic Journal: Economic Policy*, *3*(2), 97–129.

Helgelendingen. (2000, June 7). Strengere miljøkrav til aluminiumverk. Retrieved March 8, 2019, from https://www.helg.no/lokale-nyheter/strengere-miljokravil-aluminiumverk/s/1-63-283357

Henderson, J. V. (1996). Effects of air quality regulation. *The American Economic Review*, *86*(4), 789–813.

Hervik, A. & Rye, M. (2010). Drøfting av differensiert arbeidsgiveravgift som virkemiddel for utvikling av kompetansebaserte arbeidsplasser. Retrieved May 20, 2019, from https://www.regjeringen.no/globalassets/upload/krd/nou/kompetansearbeidsplasser/differensiert_arbeidsgiveravgift.pdf

Hydro. (n.d.-a). Hydro aluminium as årdal. Retrieved May 19, 2019, from https://www.hydro.com/en/about-hydro/hydro-worldwide/europe/norway/ardal/hydro-aluminium-as-ardal/

Hydro. (n.d.-b). Hydro aluminium as høyanger. Retrieved May 19, 2019, from https://www.hydro.com/en/about-hydro/hydro-worldwide/europe/norway/hoyanger/hydro-aluminium-as-hoyanger/

Hydro. (n.d.-c). Karmøy primary production. Retrieved May 19, 2019, from https://www.hydro.com/en/about-hydro/hydro-worldwide/europe/norway/karmoy/karmoy-primary-production/

Hydro. (n.d.-d). Sunndal. Retrieved May 19, 2019, from https://www.hydro.com/en/about-hydro/hydro-worldwide/europe/norway/sunndal/

Hydro. (2007, September 19). Kraftig reduksjon i utslipp av klimagasser fra norske aluminiumverk. Retrieved May 20, 2019, from https://www.hydro.com/no-NO/media/news/2007/kraftig-reduksjon-i-utslipp-av-klimagasser-fra-norske-aluminiumverk/

Innovation Norway. (n.d.). Manufacturing and process industries: Products with the worlds lowest carbon footprint. Retrieved April 1, 2019, from https://www.innovasjonnorge.no/en/start-page/invest-in-norway/industries/manufacturing-and-process-industries/

Jæger, K. J. & Budalen, A. (2019, January 26). Norske arbeidsplasser kan stå i fare hvis strømsjokket vedvarer. Retrieved May 14, 2019, from https://www.nrk.no/nordland/norske-arbeidsplasser-kan-sta-i-fare-hvis-stromsjokket-vedvarer-1.14399185

Jaffe, A. B., Newell, R. G., & Stavins, R. N. (2002). Environmental policy and technological change. *Environmental and Resource Economics*, *22*, 41–69.

Johansen, R. A. (2008). Sluttrapport - omstillingen i årdal. *Institutt for Organisasjonsutvikling*.

Kasa, S. (1999). Social and political barriers to green tax reform: The case of co_2-taxes in norway. *CICERO Policy Note*, (5).

Kasa, S. (2000). Policy networks as barriers to green tax reform: The case of co_2-taxes in norway. *Environmental Politics*, *9*(4), 104–122.

Kjensteberg, A. & Behringer, G. H. (2019, January 28). Arbeidsgiveravgift. Retrieved May 21, 2019, from https://verdtavite.kpmg.no/arbeidsgiveravgift/

Kolstad, C. J. (2011). *Intermediate environmental economics: International edition*. Oxford University Press.

Kvande, H. & Drabløs, P. (2014). The aluminium smelting process and innovative alternative technologies. *Journal of Occupational and Environmental Medicine*, (56(5 Suppl)), S23–S32.

Liu, M., Shadbegian, R., & Zhang, B. (2017). Does environmental regulation affect labor demand in china? evidence from the textile printing and dyeing industry. *Journal of Environmental Economics and Management*, *86*, 277–294.

Mabry, R. & Sharplin, A. (1986). Does more technology create unemployment? *Cato Institute Policy Analysis*, *68*.

Magnussen, I. H., Spilde, D., & Killingland, M. (2011). *Energibruk, energibruk i fastlandsnorge* (No. 9/2011). The Norwegian Water Resources and Energy Directorate. Oslo, Norway.

Manning, A. (2004). We can work it out: The impact of technological change on the demand for low-skill workers. *Scottish Journal of Political Economy*, *51*(5), 581–609.

Miljøstatus. (2018a, December 19). Klimagassutslipp fra industri. Retrieved January 21, 2019, from http://www.miljostatus.no/tema/klima/norske-klimagassutslipp/klimagassutslipp-industri/

Miljøstatus. (2018b). Norske utslipp av klimagasser. Retrieved May 13, 2019, from https://www.miljostatus.no/kilder-klimagassutslipp

Milliman, S. R. & Prince, R. (1989). Firm incentives to promote technological in pollution control. *Journal of Environmental Economics and Management*, *17*, 247–265.

Ministry of Education and Research. (2018). Nnou 2018: 2. fremtidige kompetansebehov i— kunnskapsgrunnlaget. *Norwegian Official Research Papers*, (2).

Ministry of Trade, Industry, and Fisheries. (2000, May 18). Norsk næringsvirksomhet - Metallindustrien. Retrieved January 24, 2019, from https://www.regjeringen.no/no/dokumenter/norsk-naringsvirksomhet---metallindustri/id87574/?fbclid=IwAR06cqPnZu5DU2eKrEsWjkCde4Ca4kNK0TwhHeCZMYqjAcj3mQ9jn6sCdfU

MoCE. (1981, March 13). Pollution control act. Retrieved February 4, 2019, from https://www.regjeringen.no/en/dokumenter/pollution-control-act/id171893/. (Ministry of Climate and the Environment)

MoCE. (2000). Nou 2000:1 et kvotesystem for klimagasser. *Norwegian Official Research Papers*, (1). (Ministry of Climate and the Environment)

MoCE. (2001, June 22). St.meld. nr. 54 (2000-2001): Norsk klimapolitikk. Retrieved March 6, 2019, from https://www.regjeringen.no/no/dokumenter/stmeld-nr-54-2000-2001-/id195302/sec1. (Ministry of Climate and the Environment)

MoCE. (2018, August 23). Pollution control act. Retrieved May 13, 2019, from https://www.regjeringen.no/en/topics/climate-and-environment/climate/innsiktsartikler-klima/norwegian-carbon-credit-procurement-program/id2415405/. (Ministry of Climate and the Environment)

Moen, S. E. (2007). Innovation and production in the norwegian aluminium industry. (Centre for Technology, Innovation and Culture (TIK), University of Oslo (UiO))

Moosa, I. A. (1997). On the costs of inflation and unemployment. *Journal of Post Keynesian Economics*, *19*(4), 651–666. Copyright - Copyright M. E. Sharpe Inc. Summer 1997; Last updated - 2014-05-20; SubjectsTermNotLitGenreText - US. Retrieved May 13, 2019, from https://search.proquest.com/docview/219100863?accountid=37265

Murray, M. P. (2006). Avoiding invalid instruments and coping with weak instruments. *Journal of Economic Perspectives*, *20*(4), 111–132.

NEA. (n.d.). About us. Retrieved May 13, 2019, from https://tema.miljodirektoratet.no/en/About-us/. (Norwegian Environment Agency)

NEA. (2000, May 4). Strengere miljøkrav til aluminiumsverk. Retrieved May 13, 2019, from http://test-old.miljodirektoratet.no/no/Nyheter/Nyheter/Old-klif/2000/Mai_2000/Strengere_miljokrav_til_aluminiumsverk/. (Norwegian Environment Agency)

NEA & SSB. (2013, March 15). Utslipp av klimagasser 1990 – 2010: Trender og drivkrefter. Retrieved February 14, 2019, from http://tema.miljodirektoratet.no/no/Nyheter/Nyheter/Old-klif/2013/Mars-2013/Trender_og_drivkrefter_bak_klimagassutslippene_/. (Norwegian Environment Agency Statistics Norway)

Nilsen, A. (n.d.). Nsds kommunedatabase. Retrieved May 31, 2019, from https://trygg.nsd.uib.no/nsd/nsdnytt/10-3/nsds_kommunedatabase.html. (Norwegian Centre for Research Data)

Norsk Telegrambyrå. (2003, August 15). Masseledighet truer tre hydro-samfunn. Retrieved May 14, 2019, from https://www.aftenbladet.no/okonomi/i/KW0le/Masseledighet-truer-tre-Hydro-samfunn. (Norsk Telegrambyrå)

Norsk Telegrambyrå. (2008, November 26). 450 hydro-ansatte må gå. Retrieved May 14, 2019, from https://www.dagbladet.no/nyheter/450-hydro-ansatte-ma-ga/65353942. (Norsk Telegrambyrå)

Ntb. (2014, November 12). Kina og usa fastsetter mal for co2 kutt. Retrieved May 14, 2019, from https://www.aftenposten.no/verden/i/LQ5R/Kina-og-USA-fastsetter-mal-for-CO2-kutt

OECD. (n.d.). Employment rate. Retrieved April 10, 2019, from https://data.oecd.org/emp/employment-rate.htm

Olsen, Y. (1996). *Hjørnesteinen: Elkem aluminium lista 25 år 1971-1996*. Elkem aluminium Lista.

Popp, D., Newell, R. G., & Jaffe, A. B. (2010). Energy, the environment and technological change. In B. H. Hall & N. Rosenberg (Eds.), *Economics of innovation* (Chap. 21, Vol. 2, pp. 874–937). Elsevier.

Porter, M. E. & van der Linde, C. (1995). Toward a new conception of the environment-competitiveness relationship. *The Journal of Economic Perspectives*, *9*(4), 97–118.

Regjeringen. (2000, December 11). Norsk næringsvirksomhet - metallindustrien. Retrieved January 21, 2019, from https://www.regjeringen.no/no/dokumenter/norsk-naringsvirksomhet---metallindustri/id87574/

Reitan, M. (1998). Ecological modernisation and 'realpolitik': Ideas, interests and institutions. *Environmental Politics*, *7*(2), 1–26.

Rubin, D. B. (2008). For objective causal inference, design trumps analysis. *Annals of Applied Statistics*, *2*(3), 808–840.

Schlenker, W. & Walker, R. (2011). Airport, air pollution and contemporaneous health. *NBER Working Paper Series*, (17684).

Sør-Norge aluminium A/S 1962-1987. (1987). Sør-Norge aluminium. Husnes: Sør-Norge aluminium A/S.

SSB. (n.d.). 08517: Arbeidsledige, etter kjønn og alder 1972 - 2018. Retrieved April 10, 2019, from https://www.ssb.no/statbank/table/08517. (Statistics Norway)

SSB. (2010, October 11). Prisutviklingen i metallindustrien i perioden 2000-2010: Høye priser på aluminium og nikkel. Retrieved February 27, 2019, from https://www.ssb.no/priser-og-prisindekser/artikler-og-publikasjoner/hoye-priser-paa-aluminium-og-nikkel. (Statistics Norway)

SSB. (2014, March 14). Fra bedrift til virksomhet. Retrieved May 12, 2019, from https://www.ssb.no/virksomheter-foretak-og-regnskap/artikler-og-publikasjoner/fra-bedrift-til-virksomhet. (Statistics Norway)

SSB. (2017, October 19). Viktig arbeidsgiver i distriktene. Retrieved May 15, 2019, from https://www.ssb.no/energi-og-industri/artikler-og-publikasjoner/viktig-arbeidsgiver-i-distriktene. (Statistics Norway)

SSB. (2019a). Arbeidskraftundersøkelsen. Retrieved April 9, 2019, from https://www.ssb.no/en/arbeid-og-lonn/statistikker/aku. (Statistics Norway)

SSB. (2019b, March 11). Producer price index—07201: Spot prices metals 1990 - 2018. Retrieved May 10, 2019, from https://www.ssb.no/en/statbank/table/07201/. (Statistics Norway)

Staiger, D. & Stock, J. H. (1997). Instrumental variables regression with weak instruments. *Econometrica*, *65*(3), 557–586.

Store Norske Leksikon. (2015, January 2). Skatt. Retrieved May 20, 2019, from https://snl.no/skatt#-Skattereformer

Store Norske Leksikon. (2017, May 9). Arbeidsgiveravgift. Retrieved April 3, 2019, from https://snl.no/arbeidsgiveravgift

Store Norske Leksikon. (2018a, February 12). Industri. Retrieved March 31, 2019, from https://snl.no/industri

Store Norske Leksikon. (2018b, December 12). Mosjøen aluminiumverk. Retrieved May 19, 2019, from https://snl.no/Mosj%5C%C3%5C%B8en_Aluminiumverk

The Federation of Norwegian Industries. (n.d.). Om aluminiumbransjen. Retrieved February 27, 2019, from https://www.norskindustri.no/bransjer/aluminium/om-aluminiumsbransjen/

Totten, G. & MacKenzie, D. (2003). *Handbook of aluminum: Volume 2: Alloy production and materials manufacturing*. Handbook of aluminum. CRC Press.

Vevatne, J., Lindhjem, H., Eskeland, G., Haugland, T., & Gullberg, A. (2004). Climate policy instruments and the competitiveness of norwegian industry. (8). (A work by CICERO and ECON Analyse, financed by the Ministry of Climate and the Environment)

Walker, R. (2011). Environmental regulation and labor reallocation: Evidence from the clean air act. *American Economic Review:Papers Proceedings*, *101*(3), 442–447.

Walker, R. (2013). The transitional costs of sectoral reallocation: Evidence from the clean air act and the workforce. *The Quarterly Journal of Economics*, 1787–1835.

Wooldridge, J. M. (2016). *Introductory econometrics*. Cengage Learning.

A Appendix

A.1 Trend Plots across Sample Period

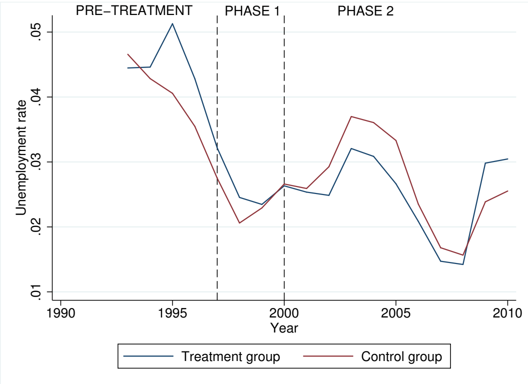

Figure 6: Trends in Unemployment Rate

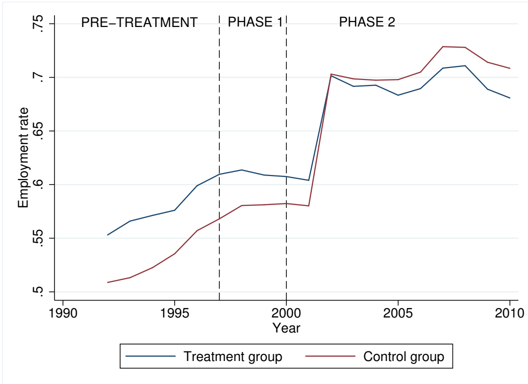

Figure 7: Trends in Employment Rate

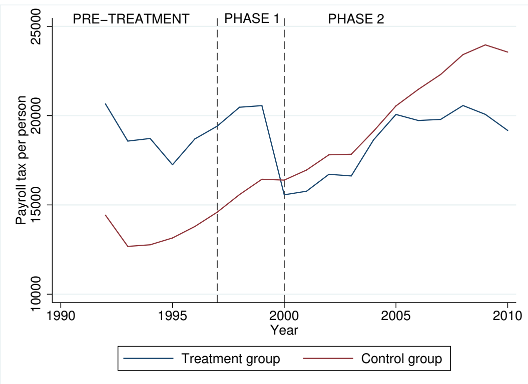

Figure 8: Trends in Payroll Tax per Person in the Working Age Population

A.2 Description of Control Variables

Table 8: List of Characteristics and Outcome Variables

	Description	Formula
A. Municipality Characteristics		
Children	1992-2005: Population between 0-15 years; 2006-2013: Population between 0-14 years.	
Youth	Population aged 16-24.	
Elder	Elder population aged 75+.	
Adult population	Population over 16 years.	
Working age population	Population in working age—aged 15-74.	Population - Children - Elder
Employed	Employed persons.	
Coastline	A dummy equal to one if a municipality has a coastline, and zero otherwise.	
B. Outcome Variables		
Unemployment Rate		
Employment Rate	The employed as a ratio to the labor force force and those outside the labor force.	Employed / Working age population
Payroll Tax	Payroll tax as a ratio to each person in the working age population.	Payroll tax / Working age population
C. Labor Market Characteristics		
Population Density	Population per square kilometer.	Population / area
Share Working Age Population	The working age population as a ratio to the population of the municipality.	Working age population / Population
Share of Women	Share of women in the population.	Women / Population
Share of women employed	Share of women from employed persons.	Employed women / Employed
Share of Youth	Share of people aged between 16 and 24 in the working age population.	Youth / Working age population
Share of Children	Share of children in the population.	Children / Population
Share of Elder	Share of elder in the population.	Elder / Population
Elementary School	Ref. Notes	Elementary education / Adult population
High School	"	High school education / Adult population
Higher Education Short	"	Higher education short / Adult population
Higher Education Long	"	Higher education long / Adult population

Notes: All variables are collected from 1992 to 2013, except for the unemployment rate, which was only available from 1993 and throughout the sample period. All variables are collected at municipality–level. Elementary school, High School, Higher Education Long and Higher Education Short are variables representing the number of people aged over 15 who have elementary school, high school, higher education below four years and higher education over four years as their highest obtained educational attainment level level. They are constructed as shares of the adult population.

A.3 Robustness Tests

A.3.1 Validation of Differential Evolution

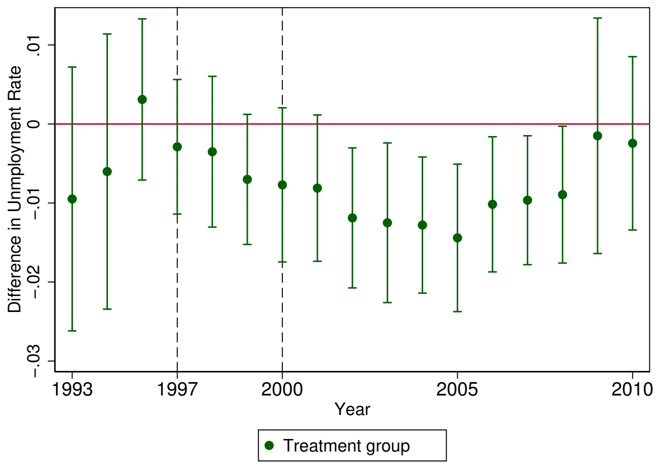

Figure 9: Differential Evolution of the Unemployment Rate

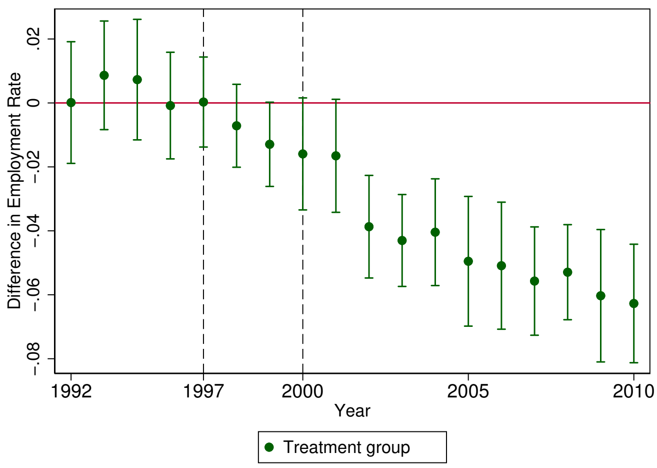

Figure 10: Differential Evolution of the Employment Rate

Notes: The figures plot the coefficients and the 5%-confidence interval (green lines) including municipality and year fixed effects, explained in 6.1.1. The horizontal line represents the control mean of the control group in each year, while the dashed horizontal indicated the agreement in 1997 and the stringency of the agreement in 2000. The year before the 97-agreement, 1996, is omitted from the plot.

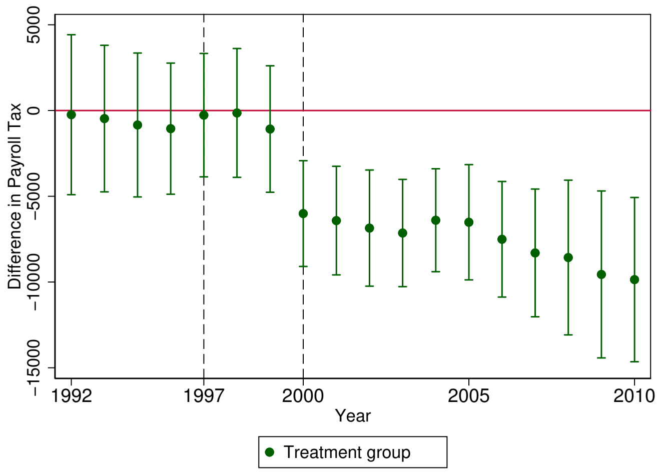

Figure 11: Differential Evolution of the Payroll Tax per person

Notes: The figure plots the coefficients and the 5%-confidence interval (green lines) including municipality and year fixed effects, explained in 6.1.1. The horizontal line represents the control mean of the control group in each year, while the dashed horizontal indicated the agreement in 1997 and the stringency of the agreement in 2000. The year before the 97-agreement, 1996, is omitted from the plot.

A.3.2 Effects on Working Age Population and Emigration

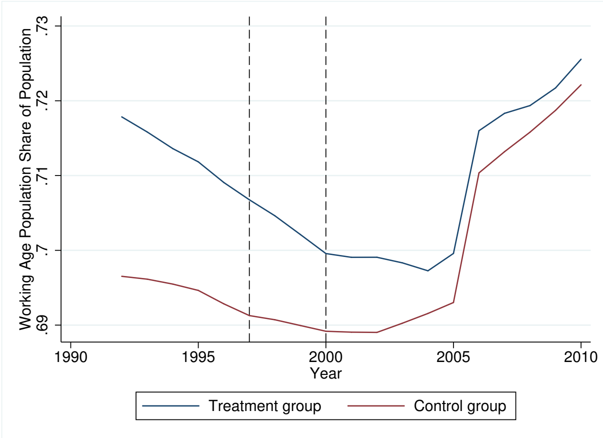

Figure 12: Trends in Working Age Population as a Share of Total Population

Table 9: DiD Estimates for Working Age Population as a Share of Total Population

	(1) Mean pre-treatment	(2) Baseline results	(3) Including controls
Working Age Population	0.714		
	(0.00389)		
Treat · Phase1		0.000450	-0.000150
		(0.00152)	(0.00139)
Treat · Phase2		-0.00363**	-0.00272**
		(0.00147)	(0.00130)
Observations	35	2869	2869
Adjusted R^2		0.819	0.828
Fixed Effects		Yes	Yes
Control Variables		No	Yes

Notes: Estimations are based on OLS of Equation 1 presented in Section 5. There are four control variables: Share of women in the workforce, share of persons with high school as their highest educational attainment level, share of persons with elementary school as their highest educational attainment level, and the aluminum price. The sample consists of 151 municipalities over the sample period 1992-2010. Phase 1 is the years 1997-1999, and Phase 2 is the years 2000-2010. Column (1) depicts the mean values in the pre-treatment period in the treatment group. Column (2) shows the baseline results from equation 1, including year and municipality fixed effects. Column (3) depicts the estimates from equation 1, including year and municipality fixed effects, as well as control variables. Standard errors are in parentheses and robust to heteroskedasticity.

* p<0.10, ** p<0.05, *** p<0.01

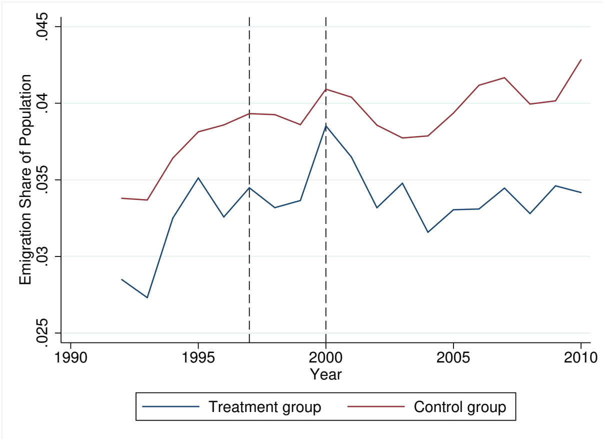

Figure 13: Trends in Emigration as a Share of Total Population

Table 10: DiD Estimates for Emigration as a Share of Total Population

	(1) Mean pre-treatment	(2) Baseline results	(3) Including controls
Emigrants	0.0312		
	(0.00188)		
Treat · Phase1		0.00154	0.00110
		(0.00161)	(0.00153)
Treat · Phase2		0.00104	0.00132
		(0.00105)	(0.000989)
Observations	35	2869	2869
Adjusted R^2		0.659	0.698
FixedEffects		Yes	Yes
ControlVariables		No	Yes

Notes: Estimations are based on OLS of Equation 1 presented in Section 5. There are four control variables: Share of women in the workforce, share of persons with high school as their highest educational attainment level, share of persons with elementary school as their highest educational attainment level, and the aluminum price. The sample consists of 151 municipalities over the sample period 1992-2010. Phase 1 is the years 1997-1999, and Phase 2 is the years 2000-2010. Column (1) depicts the mean values in the pre-treatment period in the treatment group. Column (2) shows the baseline results from equation 1, including year and municipality fixed effects. Column (3) depicts the estimates from equation 1, including year and municipality fixed effects, as well as control variables. Standard errors are in parentheses and robust to heteroskedasticity.

* $p<0.10$, ** $p<0.05$, *** $p<0.01$

A.3.3 Placebo Treatment Period

Table 11: Placebo test

	(1) Employment rate	(2) Unemployment rate	(3) Payroll tax
Treat · 1995	0.0282***	0.00949**	4192.1***
	(0.00748)	(0.00411)	(1559.5)
Treat · 1996	0.0292***	0.00638	5280.2***
	(0.00725)	(0.00412)	(1563.3)
Treat · Phase1	0.0173***	0.000309	2875.3**
	(0.00520)	(0.00253)	(1138.2)
Treat · Phase2	-0.00785	-0.00555**	-1316.5
	(0.00552)	(0.00228)	(842.2)
Observations	2869	2718	2805
Adjusted R^2	0.946	0.751	0.899
Fixed Effects	Yes	Yes	Yes
Control Variables	Yes	Yes	Yes

Notes: Estimations are based on OLS of Equation 1 presented in Section 5, expanded with interaction terms between treatment and the years 1995 and 1996. There are four control variables: Share of women in the workforce, share of persons with high school as their highest educational attainment level, share of persons with elementary school as their highest educational attainment level, and the aluminum price. The sample consists of 151 municipalities over the sample period 1992-2010. Phase 1 is the years 1997-1999, and Phase 2 is the years 2000-2010. Column (1) depicts the mean values in the pre-treatment period in the treatment group. Column (2) shows the baseline results from equation 1, including year and municipality fixed effects. Column (3) depicts the estimates from equation 1, including year and municipality fixed effects, as well as control variables. Standard errors are in parentheses and robust to heteroskedasticity.

* $p<0.10$, ** $p<0.05$, *** $p<0.01$

A.3.4 Omitting municipalities singly from the Treatment Group

Table 12: DiD Estimates for Outcomes of Labor Demand without Karmøy

	(1) Mean pre-treatment	(2) Baseline results	(3) Including controls
A. Dependent variable: Unemployment Rate			
	0.0435		
Treat · Phase1		-0.00140	-0.00137
		(0.00269)	(0.00267)
Treat · Phase2		-0.00798***	-0.00809***
		(0.00236)	(0.00234)
Observations	24	2700	2700
Adjusted R^2		0.746	0.750
B. Dependent variable: Employment Rate			
	0.578		
Treat · Phase1		0.0142**	0.0120**
		(0.00592)	(0.00537)
Treat · Phase2		-0.0158**	-0.0125**
		(0.00662)	(0.00587)
Observations	30	2850	2850
Adjusted R^2		0.937	0.945
C. Dependent variable: Payroll Tax			
	19271.5		
Treat · Phase1		2185.4*	2456.7**
		(1204.7)	(1222.2)
Treat · Phase2		-1803.3**	-2272.4**
		(859.8)	(891.2)
Observations	27	2786	2786
Adjusted R^2		0.889	0.899
Fixed Effects		Yes	Yes
Control Variables		No	Yes

Notes: Estimations are based on OLS of Equation 1 presented in Section 5. There are four control variables: Share of women in the workforce, share of persons with high school as their highest educational attainment level, share of persons with elementary school as their highest educational attainment level, and the aluminum price. The sample consists of 151 municipalities over the sample period 1992-2010. Phase 1 is the years 1997-1999, and Phase 2 is the years 2000-2010. Column (1) depicts the mean values in the pre-treatment period in the treatment group. Column (2) shows the baseline results from equation 1, including year and municipality fixed effects. Column (3) depicts the estimates from equation 1, including year and municipality fixed effects, as well as control variables. Standard errors are in parentheses and robust to heteroskedasticity.

* p<0.10, ** p<0.05, *** p<0.01

Table 13: DiD Estimates for Outcomes of Labor Demand without Årdal

	(1) Mean pre-treatment	(2) Baseline results	(3) Including controls
A. Dependent variable: Unemployment Rate			
	0.0436		
Treat · Phase1		-0.000240	-0.000238
		(0.00252)	(0.00249)
Treat · Phase2		-0.00488**	-0.00495**
		(0.00222)	(0.00219)
Observations	24	2700	2700
Adjusted R^2		0.748	0.752
B. Dependent variable: Employment Rate			
	0.5638		
Treat · Phase1		0.0135***	0.0122**
		(0.00518)	(0.00481)
Treat · Phase2		-0.0114**	-0.00945**
		(0.00516)	(0.00463)
Observations	30	2850	2850
Adjusted R^2		0.939	0.947
C. Dependent variable: Payroll Tax			
	17758.1		
Treat · Phase1		1607.4*	1778.6*
		(953.3)	(977.7)
Treat · Phase2		-1093.1	-1432.5**
		(682.0)	(709.9)
Observations	27	2786	2786
Adjusted R^2		0.891	0.901
Fixed Effects		Yes	Yes
Control Variables		No	Yes

Notes: Estimations are based on OLS of Equation 1 presented in Section 5. There are four control variables: Share of women in the workforce, share of persons with high school as their highest educational attainment level, share of persons with elementary school as their highest educational attainment level, and the aluminum price. The sample consists of 151 municipalities over the sample period 1992-2010. Phase 1 is the years 1997-1999, and Phase 2 is the years 2000-2010. Column (1) depicts the mean values in the pre-treatment period in the treatment group. Column (2) shows the baseline results from equation 1, including year and municipality fixed effects. Column (3) depicts the estimates from equation 1, including year and municipality fixed effects, as well as control variables. Standard errors are in parentheses and robust to heteroskedasticity.

* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table 14: DiD Estimates for Outcomes of Labor Demand without Høyanger

	(1) Mean pre-treatment	(2) Baseline results	(3) Including controls
A. Dependent variable: Unemployment Rate			
	0.0496		
Treat · Phase1		-0.00104	-0.00102
		(0.00249)	(0.00246)
Treat · Phase2		-0.00658***	-0.00659***
		(0.00225)	(0.00222)
Observations	24	2700	2700
Adjusted R^2		0.750	0.754
B. Dependent variable: Employment Rate			
	0.564		
Treat · Phase1		0.0111**	0.00970**
		(0.00510)	(0.00474)
Treat · Phase2		-0.0164***	-0.0141***
		(0.00572)	(0.00518)
Observations	30	2850	2850
Adjusted R^2		0.938	0.947
C. Dependent variable: Payroll Tax			
	18288.9		
Treat · Phase1		1579.2	1837.7
		(1133.8)	(1170.5)
Treat · Phase2		-1861.0**	-2146.4***
		(784.4)	(819.4)
Observations	27	2786	2786
Adjusted R^2		0.890	0.900
Fixed Effects		Yes	Yes
Control Variables		No	Yes

Notes: Estimations are based on OLS of Equation 1 presented in Section 5. There are four control variables: Share of women in the workforce, share of persons with high school as their highest educational attainment level, share of persons with elementary school as their highest educational attainment level, and the aluminum price. The sample consists of 151 municipalities over the sample period 1992-2010. Phase 1 is the years 1997-1999, and Phase 2 is the years 2000-2010. Column (1) depicts the mean values in the pre-treatment period in the treatment group. Column (2) shows the baseline results from equation 1, including year and municipality fixed effects. Column (3) depicts the estimates from equation 1, including year and municipality fixed effects, as well as control variables. Standard errors are in parentheses and robust to heteroskedasticity.

* p<0.10, ** p<0.05, *** p<0.01

Table 15: DiD Estimates for Outcomes of Labor Demand without Sunndal

	(1) Mean pre-treatment	(2) Baseline results	(3) Including controls
A. Dependent variable: Unemployment Rate			
	0.0437		
Treat · Phase1		0.000992	0.00101
		(0.00252)	(0.00249)
Treat · Phase2		-0.00499**	-0.00511**
		(0.00231)	(0.00229)
Observations	24	2700	2700
Adjusted R^2		0.747	0.751
B. Dependent variable: Employment Rate			
	0.578		
Treat · Phase1		0.0125**	0.0113**
		(0.00582)	(0.00531)
Treat · Phase2		-0.0188***	-0.0158***
		(0.00639)	(0.00572)
Observations	30	2850	2850
Adjusted R^2		0.937	0.946
C. Dependent variable: Payroll Tax			
	18338.9		
Treat · Phase1		1639.5	1763.4
		(1069.7)	(1102.4)
Treat · Phase2		-1519.7*	-1877.6**
		(780.8)	(815.2)
Observations	27	2786	2786
Adjusted R^2		0.890	0.900
Fixed Effects		Yes	Yes
Control Variables		No	Yes

Notes: Estimations are based on OLS of Equation 1 presented in Section 5. There are four control variables: Share of women in the workforce, share of persons with high school as their highest educational attainment level, share of persons with elementary school as their highest educational attainment level, and the aluminum price. The sample consists of 151 municipalities over the sample period 1992-2010. Phase 1 is the years 1997-1999, and Phase 2 is the years 2000-2010. Column (1) depicts the mean values in the pre-treatment period in the treatment group. Column (2) shows the baseline results from equation 1, including year and municipality fixed effects. Column (3) depicts the estimates from equation 1, including year and municipality fixed effects, as well as control variables. Standard errors are in parentheses and robust to heteroskedasticity.

* p<0.10, ** p<0.05, *** p<0.01

Table 16: DiD Estimates for Outcomes of Labor Demand without Farsund

	(1) Mean pre-treatment	(2) Baseline results	(3) Including controls
A. Dependent variable: Unemployment Rate			
	0.0445		
Treat · Phase1		-0.00209	-0.00205
		(0.00258)	(0.00256)
Treat · Phase2		-0.00835***	-0.00843***
		(0.00224)	(0.00220)
Observations	24	2700	2700
Adjusted R^2		0.746	0.750
B. Dependent variable: Employment Rate			
	0.580		
Treat · Phase1		0.0148***	0.0129**
		(0.00569)	(0.00515)
Treat · Phase2		-0.0166**	-0.0131**
		(0.00648)	(0.00576)
Observations	30	2850	2850
Adjusted R^2		0.937	0.946
C. Dependent variable: Payroll Tax			
	18590.3		
Treat · Phase1		1501.4	1714.0
		(1217.4)	(1245.2)
Treat · Phase2		-2367.3***	-2725.9***
		(835.7)	(871.9)
Observations	27	2786	2786
Adjusted R^2		0.889	0.899
Fixed Effects		Yes	Yes
Control Variables		No	Yes

Notes: Estimations are based on OLS of Equation 1 presented in Section 5. There are four control variables: Share of women in the workforce, share of persons with high school as their highest educational attainment level, share of persons with elementary school as their highest educational attainment level, and the aluminum price. The sample consists of 151 municipalities over the sample period 1992-2010. Phase 1 is the years 1997-1999, and Phase 2 is the years 2000-2010. Column (1) depicts the mean values in the pre-treatment period in the treatment group. Column (2) shows the baseline results from equation 1, including year and municipality fixed effects. Column (3) depicts the estimates from equation 1, including year and municipality fixed effects, as well as control variables. Standard errors are in parentheses and robust to heteroskedasticity.

* p<0.10, ** p<0.05, *** p<0.01

Table 17: DiD Estimates for Outcomes of Labor Demand without Vefsn

	(1) Mean pre-treatment	(2) Baseline results	(3) Including controls
A. Dependent variable: Unemployment Rate			
	0.0495		
Treat · Phase1		-0.00294	-0.00288
		(0.00240)	(0.00238)
Treat · Phase2		-0.00853***	-0.00857***
		(0.00214)	(0.00212)
Observations	24	2700	2700
Adjusted R^2		0.752	0.756
B. Dependent variable: Employment Rate			
	0.575		
Treat · Phase1		0.0139**	0.0120**
		(0.00578)	(0.00531)
Treat · Phase2		-0.0156**	-0.0137**
		(0.00641)	(0.00575)
Observations	30	2850	2850
Adjusted R^2		0.937	0.946
C. Dependent variable: Payroll Tax			
	19398.8		
Treat · Phase1		1688.1	1854.2
		(1191.3)	(1219.8)
Treat · Phase2		-2475.3***	-2841.7***
		(787.3)	(820.3)
Observations	30	2789	2789
Adjusted R^2		0.889	0.899
Fixed Effects		Yes	Yes
Control Variables		No	Yes

Notes: Estimations are based on OLS of Equation 1 presented in Section 5. There are four control variables: Share of women in the workforce, share of persons with high school as their highest educational attainment level, share of persons with elementary school as their highest educational attainment level, and the aluminum price. The sample consists of 151 municipalities over the sample period 1992-2010. Phase 1 is the years 1997-1999, and Phase 2 is the years 2000-2010. Column (1) depicts the mean values in the pre-treatment period in the treatment group. Column (2) shows the baseline results from equation 1, including year and municipality fixed effects. Column (3) depicts the estimates from equation 1, including year and municipality fixed effects, as well as control variables. Standard errors are in parentheses and robust to heteroskedasticity.

* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table 18: DiD Estimates for Outcomes of Labor Demand without Kvinnherad

	(1) Mean pre-treatment	(2) Baseline results	(3) Including controls
A. Dependent variable: Unemployment Rate			
	0.0461		
Treat · Phase1		-0.00248	-0.00245
		(0.00271)	(0.00269)
Treat · Phase2		-0.00812***	-0.00816***
		(0.00243)	(0.00240)
Observations	24	2700	2700
Adjusted R^2		0.746	0.750
B. Dependent variable: Employment Rate			
	0.573		
Treat · Phase1		0.0160***	0.0143***
		(0.00579)	(0.00528)
Treat · Phase2		-0.0156**	-0.0129**
		(0.00657)	(0.00585)
Observations	30	2850	2850
Adjusted R^2		0.937	0.946
C. Dependent variable: Payroll Tax			
	19391.7		
Treat · Phase1		2282.5*	2449.4**
		(1198.7)	(1232.9)
Treat · Phase2		-1872.4**	-2208.6**
		(858.6)	(893.4)
Observations	27	2786	2786
Adjusted R^2		0.889	0.899
Fixed Effects		Yes	Yes
Control Variables		No	Yes

Notes: Estimations are based on OLS of Equation 1 presented in Section 5. There are four control variables: Share of women in the workforce, share of persons with high school as their highest educational attainment level, share of persons with elementary school as their highest educational attainment level, and the aluminum price. The sample consists of 151 municipalities over the sample period 1992-2010. Phase 1 is the years 1997-1999, and Phase 2 is the years 2000-2010. Column (1) depicts the mean values in the pre-treatment period in the treatment group. Column (2) shows the baseline results from equation 1, including year and municipality fixed effects. Column (3) depicts the estimates from equation 1, including year and municipality fixed effects, as well as control variables. Standard errors are in parentheses and robust to heteroskedasticity.

* p<0.10, ** p<0.05, *** p<0.01

A.3.5 Test of Endogeneity

Table 19: Durbin-Wu-Hausman's Test of Endogeneity

	Unemployment rate		Employment rate		Payroll tax	
	(1)	(2)	(3)	(4)	(5)	(6)
p-value residual	0.000	0.022	0.000	0.000	0.000	0.000
Control Variables	No	Yes	No	Yes	No	Yes
Municipality Linear Trend	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Each column represents the p-value of the residual of separate first-stage estimations, without and with control variables, in OLS regressions on the unemployment rate, employment rate, and payroll tax per person in the working age population.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.3.6 Extended sample

Table 20: First–Stage Estimates of Søderberg–shutdown on the 2000–restriction with Extended Sample

	(1) Søderberg off Mean pre-treatment	(2) Søderberg off Baseline	(3) Søderberg off Baseline with Controls
Søderberg off	0		
Post2000		0.439***	0.438***
		(0.0501)	(0.0496)
Observations	1359	3473	3473
Adjusted R^2		0.432	0.438
Control Variables		No	Yes
Municipality Linear Trend		Yes	Yes

Notes: Each column represents results from separate OLS regressions based on Equation 3. The dependent variable is a dummy variable equal to one in the years following the shutdowns of the Søderberg anodes in the aluminum municipalities, and zero otherwise. The sample consists of 151 municipalities over the sample period 1992-2010. The control variables included are described in Section 5.3.1. Standard errors are in parentheses and are robust to heteroskedasticity.

Table 21: OLS and 2SLS Estimates for Outcomes of Labor Demand with Extended Sample

	(1) OLS	(2) OLS	(3) 2SLS	(4) 2SLS
A. Dependent variable: Unemployment Rate				
Søderberg off	-0.00586***	0.000644	-0.00561**	0.00307
	(0.00159)	(0.00153)	(0.00279)	(0.00265)
Pre-2000 mean	0.0363	0.0363	0.0363	0.0363
Observations	3171	3171	3171	3171
B. Dependent variable: Employment Rate				
Søderberg off	0.0589***	0.00830	0.0886***	0.0232**
	(0.00398)	(0.00589)	(0.0111)	(0.0103)
Pre-2000 mean	0.538	0.538	0.538	0.538
Observations	3473	3473	3473	3473
C. Dependent variable: Payroll Tax				
Søderberg off	-627.8	-4162.9***	-3141.1***	-5736.5***
	(573.6)	(746.1)	(1150.6)	(998.1)
Pre-2000 mean	12709	12709	12709	12709
Observations	3395	3395	3395	3395
Control Variables	No	Yes	No	Yes
Municipality Linear Trend	Yes	Yes	Yes	Yes

Notes: Each column represents results from separate OLS regressions based on Equation 4. The sample consists of 151 municipalities over the sample period 1992-2013. The dependent variables are the unemployment rate in Panel A, the employment rate in Panel B and the payroll tax paid by each person in the working age population in Panel C. The control variables included are described in Section 5.3.1. The pre-2000 mean is the mean of the dependent variable before 2000. Standard errors are in parentheses and robust to heteroskedasticity.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 22: Reduced-form Estimates for Outcomes of Labor Demand with Extended Sample

	(1)	(2)
A. Dependent variable: Unemployment Rate		
Stringency in regulation 2000	-0.00248*	0.00134
	(0.00127)	(0.00114)
Pre-2000 mean	0.0363	0.0363
Observations	3171	3171
Adjusted R^2	0.002	0.234
B. Dependent variable: Employment Rate		
Stringency in regulation 2000	0.0391***	0.0102**
	(0.00452)	(0.00434)
Pre-2000 mean	0.538	0.538
Observations	3473	3473
Adjusted R^2	0.057	0.636
C. Dependent variable: Payroll Tax		
Stringency in regulation 2000	-1386.5***	-2511.5***
	(461.3)	(438.8)
Pre-2000 mean	12709	12709
Observations	3395	3395
Adjusted R^2	0.117	0.362
Control Variables	No	Yes
Municipality Linear Trend	Yes	Yes

Notes: Each column represents results from separate OLS regressions based on Equation 5. The dependent variables are the unemployment rate in Panel A, the employment rate in Panel B and the payroll tax by each person in the working age population in Panel C. The sample consists of 151 municipalities over the sample period 1992-2013. The covariates included are described in Section 5.3.1. Standard errors are in parentheses and robust to heteroskedasticity.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$