On the Effects of Energy Taxation and Emissions Trading: Evidence from German Manufacturing

Inauguraldissertation

zur Erlangung des akademischen Grades
Doctor Rerum Politicarum

an der

Fakultät für Wirtschafts- und Sozialwissenschaften
der Ruprecht-Karls-Universität Heidelberg

vorgelegt von

Benjamin Johannes Lutz

Erstgutachter: Prof. Dr. Andreas Löschel, Universität Münster
Zweitgutachter: Prof. Dr. Christian Conrad, Universität Heidelberg
Heidelberg, den 27.10.2016

Acknowledgements

I am greatly indebted to my first supervisor Andreas Löschel for his time, advice, and support - that went way beyond this thesis. Our conversations have been a great source of motivation and inspiration.

I am grateful to my second supervisor Christian Conrad who accepted to join the thesis committee without hesitation.

I thank my co-authors Florens Flues, Uta Pigorsch, Waldemar Rotfuß, and Shunsuke Managi for the fruitful collaboration.

During the last five and a half years, I have worked at the Centre for European Economic Research in Mannheim. A great deal of this thesis is based on the research I conducted there. I would like to thank my colleagues: We have laughed together and we have argued with each other - all in all, I had a splendid time. I owe great thanks to Claudio Baccianti, Kathrine von Graevenitz, Katrin Sommerfeld, Philipp Massier, Ulrich Wagner, Martin Achtnicht, Sascha Rexhäuser, Martin Kesternich, and Sebastian Voigt.

In 2015, I was given the opportunity to enjoy a research stay of four months at the University of California, Berkeley. I would like to thank Maximilian Auffhammer and Meredith Fowlie for their suggestions and the stimulating discussions.

Furthermore, I would like to thank the staff at the Research Data Centres of the Statistical Offices of the Federal States for providing data access and assistance. In particular, I owe great thanks to Michael Rößner for his patience and his support.

Finally, I am deeply grateful to my parents Ulrike and Klaus Hermann Lutz as well as the love of my life Bianca Pabst. The completion of this thesis did not only require considerable dedication from my side - but also from yours. Thank you.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 General remarks

Market-based instruments have become an important cornerstone of European climate policy. In 2005, the European Union (EU) established the worldwide largest multinational cap-and-trade system in order to limit industrial greenhouse gas emissions. Today, the EU Emissions Trading System (ETS) covers about 45 percent of the EU's greenhouse gas emissions and regulates more than 11,000 installations in 31 countries. On the national level, other market-based instruments have been implemented, too. Several European countries including Sweden, Denmark, Norway, and Germany have initiated environmental tax reforms in the last decades. Germany introduced a new electricity tax in 1999 and raised both electricity and existing fuel tax rates during the subsequent years in order to internalize the social cost of carbon.

Market-based instruments have been internationally endorsed and widely implemented in order to reduce greenhouse gas emissions from industrial sources. Nevertheless, the empirical evidence on their functioning and their causal effects on regulated firms is still scarce. This thesis contributes to two different strands of empirical literature examining market-based instruments. First, it sheds light onto the causal effects of emissions trading and energy taxation on the economic performance of regulated manufacturing firms. Exploiting official and administrative firm-level micro data from Germany, quasi-experimental approaches are employed to identify and quantify the impact of the EU ETS and the German electricity tax on the productivity, technical efficiency, and competitiveness of regulated firms. Second, it extends the literature on the price formation process on the European carbon market.

The first paper investigates the impact of the EU ETS on the productivity of German manufacturing firms. The firm-level productivity is estimated based on an empirical production function that allows for an endogenous dynamic productivity process influenced by the EU ETS. The second paper examines the effect of the EU ETS on the technical efficiency of German manufacturing firms. The third paper provides an ex-post evaluation of the German electricity tax with a focus on its effects on firm performance in the manufacturing sector. Finally, the fourth paper investigates the non-linear relationship between the EU Allowance (EUA) price and its influencing factors.

The remainder of the introduction is structured as follows. Section 1.2 gives background information on market-based instruments and the application of this intervention class in environmental policy. Section 1.3 describes recent advances in the literature on the ex-post evaluation of policy interventions. Section 1.4 summarizes the existing empirical literature on the causal effects of environmental regulation on firms. Section 1.5 provides an overview of the empirical research on the price formation process on the European carbon market. Section 1.6 outlines each of the four papers and its contribution to the literature. Finally, Section 1.7 concludes with avenues for future research.

1.2 Market-based instruments - a brief overview

Climate change poses major risks to mankind and is one of the greatest environmental challenges ever faced. In the framework of the Paris Agreement under the United Nations Framework Convention on Climate Change (UNFCCC), 191 countries signed a climate treaty that governs efforts for the achievement of a global global temperature target. The rise of the average global temperature should not exceed two degrees during the current century. Although the ratification process of the treaty is still ongoing, it is likely that the Paris Agreement will strengthen the mitigation of greenhouse gas emissions.

The need for effective and economically efficient policy instruments is strong. First, the ambitious global temperature target requires immediate and effective mitigation of greenhouse gas emissions. Second, from a welfare economic point of view, the mitigation should be achieved at the lowest possible cost. Minimizing the cost of regulation also facilitates the political feasibility of implementing mitigation measures.

The theoretical economic literature has been concerned with the market failure of negative externalities since the early 20th century. Pigou (1920) formalized the underlying principles of externalities. He suggested to implement a tax to price activities associated with negative externalities. In his theoretical framework, the optimal tax rate corresponds to the marginal external cost.

About forty years later, Coase (1960) made a seminal contribution that prepared the ground for cap-and-trade systems. He proposed that the allocation of property rights associated with the externality could be the basis for a bargain among the involved parties. Then the optimal level of pollution would be achieved through exchange without further intervention. Crocker (1966) and Dales (1968) promoted the concept of tradable pollution rights pointing out its advantages in comparison to the Pigouvian tax. Originating from an optimal level of environmental quality - or level of pollution - the regulating entity needs less information about polluters' abatement costs in order to achieve the optimal solution. In the framework of tradable pollution rights, the market mechanism will enable the formation of a price that equates the marginal external cost provided that the total amount of pollution rights corresponds to the targeted level of pollution. For a solution according to the principles of the Pigouvian tax, the regulating entity would need information on the marginal external cost in order to set the optimal tax rate. An additional favorable trait of cap-and-trade systems is that, from a theoretical perspective, the overall economic outcome is independent from the initial allocation or distribution of the pollution rights (Montgomery, 1972). In other words, no matter whether polluters receive pollution rights for free or have to buy them, the aggregate emissions target is achieved at lowest cost.

The properties of the alternative policy instruments are mostly derived in a static context relying on assumptions regarding the cost of information, uncertainty, and transaction costs. In the following decades, the theoretical literature has explored the properties of these instruments under alternative assumptions. Weitzman (1974), for instance, examined the differences between the two concepts taking into account uncertainty. Roberts and Spence (1976) investigated the properties of a hybrid system of tradable emission rights with minimum and maximum prices. Milliman and Prince (1989) analyze the dynamic efficiency of market-based instruments and command-and-control regulation showing that environmental taxation and emissions trading provide the highest incentives for regulated firms to invest in technological change. A survey of the theoretical literature on the design of interventions can be found in Cropper and Oates (1992) and Aldy, Krupnick, Newell, Parry, and Pizer (2010).

Although in theory market-based instruments are capable of achieving a socially desired environmental target at least cost, large environmental programs regulating air and water pollution were primarily build on command-and-control approaches until the late 1980s. The United States (US) adopted a pioneer role in environmental protection implementing initiatives dedicated to monitor and regulate pollution, such as the 1963 and 1970 Clean Air Act, the 1972 Clean Water Act and the 1977 Amendments to the 1970 Clean Air Act. The instruments that were implemented in the context of these programs to improve environmental quality were mainly technology-based standards. The 1978 Energy Act introduced a tax - however, it was not congruent with the Pigouvian line of thought. Instead of taxing the pollutant or a good closely related to the pollutant, the Gas Guzzler Tax focuses on technology adoption and thus applies to the sale of cars with a low mileage per gallon.

Taxes on energy use had been introduced in the US and many European countries during the first half of the 20th century. However, the incentive effect of these excise taxes on fuel use was not the impetus for their implementation. The main purpose was to generate tax revenues inter alia used to fund the infrastructure for the emerging traffic flows of motorized vehicles. In the 1970s, under the Nixon administration, a tax on sulfur dioxide was proposed to curb emissions, but these efforts were not successful (Milne, 2011).

Market-based instruments became more popular in the late 1980s and early 1990s. In 1989, the US implemented an excise tax on the sale of chlorofluorocarbons (CFCs) and other ozone-depleting chemicals in order to meet its obligations under the Montreal Protocol. In the framework of the 1990 Amendments of the 1970 Clean Air Act, the US decided to introduce the Acid Rain Program - the first large-scale cap-and-trade system to be implemented in 1995 to regulate sulfur dioxide emissions. Ellerman, Joskow, Schmalensee, Montero, and Baily (2000, p. 5) review the first years of the Acid Rain Program and come to the conclusion that it has " [...] performed well and has thereby proven that emissions trading has considerable potential in practice, [...] ". In addition, the 1990 Amendments of the 1970 Clean Air Act encouraged the implementation of market-based instruments on state level, for instance, to regulate also other local pollutants such as nitrogen oxides emissions.

At the same time, several European countries introduced environmental tax reforms in order to improve environmental quality. The reforms mainly aimed at reducing sulfur dioxide and greenhouse gas emissions (Bosquet, 2000). The enacted taxes either directly price the emission of pollutants or indirectly punish emissions by increasing the price of energy. In contrast to the early adopted fuel taxes, now, policy makers put emphasis on the incentive effects of energy and emission taxes. Germany enacted its environmental tax reform in 1999 by increasing the existing fuel taxes and introducing a new excise tax on electricity use. The EU initiated the attempt to implement a European carbon energy tax in 1992. However, the plan was abandoned in 1997, due to political reasons, such as the opinion that a European wide tax would interfere with the member states' autonomy in taxation (Ellerman, Convery, and De Perthuis, 2010).

Similarly, market-based instruments were increasingly endorsed at the international level - especially in the context of the emerging efforts of the United Nations (UN) to combat anthropogenic climate change. The UN Framework Convention on Climate Change (UNFCCC) established in 1992 was the first international treaty that recognized anthropogenic climate change and the necessity to take measures to reduce greenhouse gas emissions. It laid the ground for international negotiations on binding emission targets that were achieved in the framework of Kyoto Protocol passed in 1997. The Kyoto Protocol embraced the principles of tradable emission rights considering emissions trading a way to achieve the greenhouse gas reduction target at lowest cost and to strengthen the treaty by easing the political feasibility. It included three flexibility mechanisms that enabled parties to offset emissions, namely International Emissions Trading, the Clean Development Mechanism (CDM), and Joint Implementation (JI).

In the framework of the Kyoto Protocol, the fifteen member states of the EU committed themselves to jointly reduce greenhouse gas emissions by 8 percent in comparison to 1990 levels during the 2008-2012 Kyoto Protocol period. Soon after the Kyoto Protocol had been negotiated, the EU decided to set up its own emissions trading system among the member states in order achieve the emissions target at least cost. According to Ellerman, Convery, and De Perthuis (2010, p.16), the EU ETS was the product of two failures: "[...] the Commission's failure to win pan-European support for the introduction of a carbon tax, and the failure to insert their desired policy initiatives in the Kyoto Protocol." In 2003, the EU decided to launch the EU ETS by enacting the Emissions Trading Directive.

Furthermore, it was subsequently decided to link the EU ETS with the flexible mechanisms of the Kyoto Protocol CDM and JI. Finally, the EU ETS was launched in 2005 regulating more than 11,000 industrial installations in the back then 27 member states covering about 45 percent of their total emissions. It is the cornerstone of the EU's climate policy that has set ambitious emission reduction targets in the context of the 2020 Climate and Energy Package, the 2030 Climate and Energy Framework, and the 2050 Low-carbon Economy Roadmap. Apart from its importance for the European climate policy agenda, the EU ETS has adopted an exemplary character for other schemes that have already come into existence or will be launched in the future, such as the Regional Greenhouse Gas Initiative in the US or the Chinese emissions trading initiative.

Lately the parties of the UNFCCC have agreed on a new climate treaty succeeding the Kyoto Protocol. The Paris Agreement aims at keeping the rise of the average global temperature below two degrees during the current century. Vast emission reductions will be necessary to achieve this ambitious objective intensifying the need for effective and efficient climate policy instruments.

While the functioning of market-based instruments is well understood from a theoretical point of view, empirical evidence on the causal effects of the described interventions is still scarce. In practice, the interventions to price the externalities of greenhouse gas emissions are the outcome of political bargaining processes. Furthermore, their implementation faces issues, such as costly information, transaction costs, uncertainty, and other barriers. An ex-post evaluation of the introduced interventions is necessary to improve the design of existing as well as planned interventions.

1.3 Recent advances in the evaluation of policy interventions

Research on the evaluation of policy and programs has made significant progress during the last decades. As a result, contemporary studies focus on the identification and quantification of causal effects induced by the intervention under examination. Several methods have been developed and refined to study the impact of interventions before and after their implementation - a development that also has been pushed by the improved access to micro-level data (Imbens and Wooldridge, 2009; Angrist and Pischke (2010)).

An important factor for the evolution of this literature was the formulation of an universal concept, that could be employed to describe causal inference in very different settings (Imbens and Wooldridge, 2009): The seminal work of Rubin (1974, 1977) led to the approach that is now commonly used to describe empirical strategies for the identification of causal effects. The potential outcome framework or Rubin causal model (Holland, 1986) differentiates between the two potential outcomes that could materialize depending the participation in the intervention:

$$Y_i = \begin{cases} Y_i(0) & \text{if } D_i = 0 \\ Y_i(1) & \text{if } D_i = 1 \ , \end{cases} \tag{1.1}$$

where i indexes the subjects, Y denotes the outcome variable, D is an indicator variable that takes the value one if subject i is treated and zero otherwise. In consequence of the intervention, only one of the two potential outcomes materializes and is observable at the individual level.

There are different concepts to measure the impact of an intervention (Imbens and Wooldridge, 2009; Pearl, 2000). Most commonly used in the policy and program evaluation literature is the average treatment effect on the subjects that participated in the intervention, i.e. the average treatment effect on the treated (ATT):

$$ATT = E[Y(1)_i|D_i = 1] - E[Y(0)_i|D_i = 1]. \tag{1.2}$$

The ATT is the expected difference between the two potential outcomes conditional on the participation in the intervention. This hypothetical concept measures the average difference between the materialized - or factual - outcome and the counterfactual outcome. Only the former is realized and can be observed. This dilemma is solved by comparing the average outcome of the group of treated subjects with the average outcome of a group of similar subjects that are not treated - the control group. The selection into both groups should be independent from the outcome as well as from other factors influencing the outcome - otherwise, the expected difference between the outcomes of the two groups is not equivalent to the ATT (Angrist and Pischke, 2009).

By design, the randomized controlled trial solves this problem based on randomization: A group of randomly selected subjects participates in the intervention of interest and a disjunct group of randomly selected subjects does not participate. The comparison of the mean outcomes of the two groups yields the ATT of the policy intervention. Randomized controlled trials have been successfully implemented to study the causal effects of interventions ex-ante in different areas of economics - especially in development economics.

In addition, several methods have been developed that enable the identification of causal effects after the implementation of interventions using observational data. The adopted identification strategies rely on variation in treatment or treatment intensity, that might occur over time, over groups of subjects, or spatially. In contrast to the randomized controlled trial, the treatment variation, i.e. the selection into treatment and control group, is not random. In many cases, the variation is due to exemption schemes that exclude subjects from the intervention. The methods used for the empirical ex-post evaluation of policies address the aforementioned selection problem and other challenges that might impede the identification of the causal effect. Among the prominently used statistical and econometric tools are panel regression models, the difference-in-differences approach, the regression discontinuity and kink design, as well as instrumental variable and matching approaches.

Although the randomized controlled trial is considered the gold standard for isolating and quantifying the causal effect of a policy intervention, barriers such as high cost, ethical controversy, and a lack of acceptance among policy makers hinder the extensive use of this policy evaluation tool. This especially holds for environmental and climate policy interventions that aim at reducing industrial emissions or energy use. Therefore, this strand of literature focuses on the ex-post evaluation of already implemented interventions.

1.4 The impact of environmental regulation on firms

The recent advances in the policy and program evaluation literature have also influenced the evaluation of environmental regulation that addresses industrial pollution. The research design of empirical studies in this area increasingly emphasizes causal inference by carefully comparing groups of regulated firms with adequate groups of control firms.

Command-and-control regulation The intensive implementation of command-and-control regulation to regulate pollution of industrial installations lends plenty of variation that can be exploited to study its causal effects. A particularly fruitful setting has been provided by the 1970 Clean Air Act and its 1977 Amendment that created spatial variation in environmental regulation across the US. The key instrument of this legislation was the assignment of a county-level ambient air quality status. The nonattainment of the environmental target induced the implementation of regional technology standards.

Empirical investigations of the effects of the Clean Air Act on the behavior of plants show, that industries affected by the regulation moved their activities to counties with lower regulatory burdens (Henderson, 1996; Becker and Henderson, 2000). Greenstone (2002) provides evidence that higher regulatory burdens also decreased employment, capital stock, and output of pollution intensive industries. List, Millimet, Fredriksson, and McHone (2003) support these findings by showing that plant births, closures, and investment patterns have been adversely affected by the Clean Air Act. Hanna (2010) investigates the effect of environmental regulation on outbound foreign direct investments. The results indicate that the Clean Air Act caused regulated multinational firms to increase their foreign assets and their foreign output. Greenstone, List, and Syverson (2012) investigate the effects of the different components of the Clean Air Act on the total factor productivity of firms from the manufacturing industry. They show that air quality regulations in the US significantly decreased the total factor productivity of manufacturing plants in regulated areas.

Walker (2011) and Walker (2013) investigate the effect of the Clean Air Act on employment and labor reallocation. The evidence presented in these studies points to negative effects on employment growth, job creation rate, and a positive effect on job destruction rates. According to Walker (2013), the Clean Air Act caused large transitional costs due to the reallocation of workers.

While the policy evaluation literature provides profound evidence on the negative effect of the nonattainment status and the associated technology standards on the economic performance of regulated firms, its effect on ambient air quality is controversial. Greenstone (2004) shows that the attainment and nonattainment categories in the framework of the Clean Air Act can only explain a small fraction of the improvements of ambient sulfur dioxide concentrations.

Apart from the Clean Air Act, also other environmental command-and-control programs have been under investigation. Berman and Bui (2001) analyze the effect of command-and-control regulation on plant entry and exit exploiting variation in nitrous oxides regulation caused by the South Coast Air Quality Management District (SCAQMD). They show that there is no evidence of an adverse impact of the implemented regulation on firms in the Los Angeles area. Gray, Shadbegian, Wang, and Meral (2014) investigate the effect of multi media regulation for the pulp and paper industry on labor demand. The regulation under investigation has been implemented in the framework of the Clean Water Act and the Clean Air Act. Different kinds of technology standards have been imposed to regulate plants' air and water pollution. The results indicate that the regulation had only a limited effect on employment.

Causal effects of environmental taxation on manufacturing firms Empirical evidence allowing causal inference with regard to the effects of environmental or energy taxation on industrial installations is still scarce. Existing studies exploit exogenous variation in energy or electricity prices in order to isolate the the effect from confounding factors.

Kahn and Mansur (2013) exploit county-level spatial variation in electricity prices and labor and environmental regulation to study the impact of these factors on US manufacturing. Using industry-level data, they find that energy-intensive industries are mainly located in counties with low electricity prices.

Martin, de Preux, Wagner (2014) investigate the effect of a carbon tax - the Climate Change Levy - on manufacturing plants in the UK. The underlying identification strategy relies on variation in the tax liability. A subset of plants from energy intensive industries was granted tax reductions. The employed instrumental variable approach shows that the carbon tax had a strong negative effect on energy intensity and electricity use, while it did not affect employment, revenue, or plant exit.

Harrison, Hyman, Martin, and Nataraj (2015) exploit regional variation in environmental command-and-control regulation and coal prices in India investigating the effect of both instruments on industrial installations. They show that higher coal prices reduced coal use and improved the environmental performance of plants while the command-and-control regulation did not affect the use of coal on the intensive margin. However, the regulation increased the share of large plants that are more likely to invest in pollution control. In comparison to command-and-control regulation, higher coal prices are more effective in reducing ambient sulfur dioxide concentration.

Gerster (2015) examines the effect of electricity prices on German manufacturing plants. The German government imposed a levy on electricity use to finance the feed-in tariff system for renewable energy. The levy is not applied to electricity intensive plants creating discontinuities in the price-quantity relationship of electricity. Gerster (2015) employs a regression discontinuity design to identify the effect of electricity price on electricity use and economic performance. He shows that electricity use reacts to prices, however there is no evidence of a negative effect on gross output, exports, or employment.

The forth chapter of this thesis contributes to this literature by examining the effect of the German electricity tax on the competitiveness of manufacturing firms. The analysis of the impact of the electricity tax focuses on firms that use less than one gigawatt hour, while Gerster (2015) addresses firms that use more than ten gigawatt hours of electricity. The results of the forth chapter are in line with Gerster (2015). The empirical analysis based on the regression discontinuity design provides no evidence for a statistically significant effect of the electricity tax on revenues, value added, exports, or employment.

Causal effects of emissions trading systems on firms In order to curb sulfur dioxide and nitrous oxides, several regional emissions trading systems have been implemented in the US. In addition, the EU ETS offers the possibility to investigate the causal effects of emissions trading on regulated firms. Despite these settings akin to natural experiments, the empirical literature on the causal effects of cap-and-trade systems is still scarce.

Fowlie, Holland, and Mansur (2012) investigate the causal effect of the Californian REgional CLean Air Incentives Market (RECLAIM) on the emissions of regulated installations. RECLAIM is a regional cap-and-trade system regulating nitrogen oxides emissions of installations in southern California. The identification strategy relies on treatment variation due to the design of the regulation: only a subset of industrial installations is regulated by RECLAIM, the remaining installations are regulated based on command-and-control regulation. Fowlie, Holland, and Mansur (2012) compare the average emissions of the RECLAIM installations with the average emissions of an adequate control group that is constructed using semiparametric matching. The results suggest that installations under emissions trading reduced their emissions significantly in comparison to the installations regulated by command-and-control regulation. Furthermore, they show that emissions reductions are equally distributed over areas with different socioeconomic backgrounds resolving doubts that RECLAIM led to environmental injustice. Fowlie and Perloff (2013) empirically investigate whether the economic outcome of emissions trading is independent from the mode of allowance allocation. In the framework of RECLAIM, the timing of the permit allocation offers exogenous variation that can be exploited to identify the effect of the mode of allocation on nitrogen oxide emissions. The results indicate that the emissions reductions do not depend on how the permits were allocated to firms..

Linn (2011) analyzes the effect of the nitrogen oxides budget trading program on firm profits. The program covers the eastern US regulating emissions of plants from the electricity sector. The results indicate a negative impact of the program on expected profits of regulated firms.

Ferris, Shadbegian, and Wolverton (2014) analyze the effect of the sulfur dioxide trading program on electricity utility employment. The program has been enacted in the framework of the 1990 Clean Air Act Amendments. The effect of the trading program on employment is not statistically significant at the utility level.

Most recently, a growing literature on the evaluation of the EU ETS has emerged. The source of variation to identify the causal effects of the EU ETS is its inclusion restrictions. The European Commission granted exemption from the scheme for small emitters in order to balance transaction costs. The inclusion into the scheme is therefore a function of installation-level capacity. Within narrowly defined industries, the EU ETS leaves a subset of firms unregulated that may serve as control group provided that the research design takes into account that the capacity and thus the treatment is correlated with the outcomes of interest.

With regard to carbon dioxide emissions the empirical evidence is rather clear: Investigating firm- and plant-level data for different European countries, a number of studies show that the EU ETS significantly reduced the emissions or the emissions intensity of regulated entities (Petrick and Wagner, 2014; Wagner, Muûls, Martin, and Colmer, 2014; Klemetsen, Rosendahl, and Jakobsen, 2016; Jaraitė and Di Maria, 2016).

The empirical evidence on the effect of the EU ETS on employment points in different directions depending on the data examined. Petrick and Wagner (2014) do not find a statistically significant effect of the EU ETS on employment using firm-level data from the German manufacturing sector. The companion study by Wagner, Muûls, Martin, and Colmer (2014) exploiting plant-level data from the French manufacturing sector, however, finds a significant negative impact of the EU ETS on employment of regulated plants. Abrell, Ndoye Faye, and Zachmann (2011) do not find a statistically significant effect of the EU ETS on employment analyzing European commercial firm-level data.

Empirical evidence also suggests that the EU ETS did not have a negative effect on economic performance measured by output, exports, labor productivity or value added. Abrell, Ndoye Faye, and Zachmann (2011) are not able to reject the null hypothesis that the EU ETS had no effect on output and value added of regulated firms. Petrick and Wagner (2014) find that the EU ETS had a positive effect on the gross output of German manufacturing firms while it did not affect exports. Jaraitė and Di Maria (2016) do not find a significant effect of the EU ETS on profitability of Lithuanian firms. Klemetsen, Rosendahl, and Jakobsen, 2016 find a positive effect of the EU ETS on value added and labor productivity.

Calel and Dechezleprêtre, 2016 investigate the impact of the EU ETS on low carbon innovation by analyzing patenting activities of European firms. They find that the EU ETS increased the number of successful low carbon patent applications of regulated firms.

The second and third chapter of this thesis contribute to this literature by investigating the causal effect of the EU ETS on the economic performance of German manufacturing firms. Two different comprehensive measures for firm-level economic performance are estimated. The second chapter employs a structural production function model by Ackerberg, Caves, and Frazer (2015) and De Loecker (2013) in order to estimate the firm-level total factor productivity. The empirical analysis in the third chapter is based on the stochastic frontier approach by Aigner, Lovell, and Schmidt (1977) that is used to derive firm-level technical efficiencies. The results of both chapters suggest no statistically significant impact of the EU ETS on economic performance.

All in all, the evaluation literature predominantly provides empirical evidence for a significant positive effect of market-based instruments on the reduction of emissions and energy use, while the effect of command-and-control regulation seems to be controversial, at least with regard to parts of the Clean Air Act legislation (Greenstone, 2004). On the other hand, so far, there is little evidence of negative effects of market-based instruments on economic performance of regulated firms, while command-and-control regulation has been found to have negative effects on output, employment, and total factor productivity in several cases.

1.5 Price formation on the European allowance market

The efficient formation of the allowance price is an important prerequisite for a functioning cap-and-trade system. According to the price, regulated firms decide whether they individually abate emissions or buy allowances. For regulating entities, the price is an important signal giving information about the stringency of the scheme. Optimal decisions can only be taken by the market participants if the allowance price reflects all available information on the cost of abatement to achieve the emissions target (Hintermann, Peterson, and Rickels 2016).

During the last years, a large body of empirical literature has formed that investigates the price formation process on the market for allowances traded in the framework of the EU ETS: European Union Allowances (EUAs). The early studies employed univariate time series models in order to investigate the characteristics of the EUA price. Autoregressive - generalized autoregressive conditional heteroscedasticity (AR-GARCH) models that take into account autocorrelation and conditional heteroscedasticity have played a prominent role in this literature (Paolella and Taschini, 2008; Benz and Trück, 2009).

In addition several studies analyze the relationship between the EUA spot and derivative markets. Option pricing models have been used to examine the price formation on different EUA futures markets and their relationship to the spot market (Chevallier, Ielpo, and Mercier; 2009; Daskalakis, Psychoyios, and Markellos, 2009). These studies emphasize the importance of the futures markets for the processing of information in the price formation process. Chevallier, Le Pen, and Sévi (2011) show that the introduction of an EUA options market decreased the volatility of the EUA prices on other markets.

A further topic under investigation has been the relationship between the EUA price time series and its determinants. Empirical evidence for the first years after the implementation of the EU ETS points to a strong relationship between the EUA price and energy commodity prices, such as the prices for crude oil, gas, and coal while the influence of temperature plays only a minor role (Mansanet-Bataller, Pardo, and Valor, 2007; Alberola, Chevallier, and Chèze, 2008). Chevallier (2009) shows that also financial and macroeconomic risk factors influence the EUA price. While the results of these studies are mainly based on linear regression models, Hintermann (2010) employs a structural approach that takes into account the specific characteristics of the demand side of the European allowance market. His results support the strong influence of energy prices on the EUA price. Creti, Jouvet, Mignon (2012) study the longterm relationship of the EUA price and its determinants employing a cointegration approach. They show that according to the equilibrium relationship between the time series the EUA price has been undervalued since 2009. Aatola, Ollikainen, and Toppinen (2013) also investigate the determinants of the EUA price. Employing different linear regression models, they show that the difference between gas and coal prices, as well as electricity prices influence the EUA price. Koch, Fuss, Grosjean, and Edenhofer (2014) shed light onto abatement-related determinants of the EUA price. They find that the influence of the deployment of renewables and the use of offset credits is only moderate in comparison to the influence of economic activity.

These results are mostly based on linear regression models that consider structural breaks in the EUA price time series by analyzing subsamples. Structural breaks in the data generating process are an important feature of the EUA price time series. Chevallier (2011a, 2011b) and Peri and Baldi (2011) explore the nonlinear relationship of the EUA price time series and its fundamentals using regime-switching models. Their results corroborate the strong influence of economic activity and energy prices on the EU ETS. Koch (2014) examines the time varying relationship between the EUA price and its fundamentals by estimating a econometric model with different regimes. It allows the relationship to smoothly change across regimes. His results show, that the relationship between the EUA price and the energy prices became stronger in the years from 2008 to 2012 in comparison to the years directly after the implementation of the EU ETS.

The fifth chapter of thesis contributes to the literature on the nonlinear relationship between the EUA price and its fundamentals. A Markov regime-switching GARCH model is implemented in order to endogenously account for structural breaks in the price formation process. A low and a high volatility regime are identified showing a changing relationship between the EUA price and its fundamentals. Within this empirical framework, the most important price determinants are the stock market and energy prices.

While the former studies are based on daily or even higher aggregated price and return series, Rittler (2012) and Conrad, Rittler, and Rotfuß (2012) exploit intraday data to study the EUA price formation process. Rittler (2012) analyzes the relationship between EUA spot and futures prices, in particular, the transmission of information in the first and second conditional moments. The results indicate that the futures market takes up information first and then passes it on to the spot market. Conrad, Rittler, and Rotfuß (2012) investigate the impact of information disclosure on the EUA price formation. They show that the information on allowance allocation, cap, and future economic development are immediately processed and reflected by the EUA price. The impact of regulatory announcements on the EUA price is supported by the findings of Koch, Grosjean, Fuss, and Edenhofer (2016). They apply an event study method to examine the impact of information disclosure on allowance supply schedules as well as European climate policy targets. Mizrach and Otsubo (2014) examine the market microstructure of the largest European carbon trading platform, the European Climate Exchange. Their analysis of the realized volatility, bid-ask spreads and adverse selection costs reveals imbalances in the order book that can be exploited for a profitable trading strategy. Medina, Pardo, and Pascual (2014) also investigate the microstructure of the EUA market. They show that trading frictions measured by relative spreads, information-asymmetry risk, and market-making profits decreased during the years from 2008 to 2012.

Reboredo (2013) and Balcilar, Demirer, Hammoudeh, and Nguyen (2016) take the point of view of institutional investors and analyze the interdependency of the carbon and energy markets. They come to the conclusion that the carbon market can help to diversify portfolio risk.

Altogether, the empirical literature on the price formation on the European carbon market points to three important insights. First, economic activity, energy prices, and regulatory announcements are the most important drivers of the EUA price. Second, the price formation process is characterized by breaks that also translate into the relationship between the EUA price and its fundamentals. Third, the market for EUAs became more mature during the years from 2008 onwards.

1.6 Outline and findings of the thesis

This thesis is based on four independent research papers that contribute to the two strands of literature on market-based instruments depicted above. The second and the third chapter investigate the effects of the EU ETS on the economic performance of regulated manufacturing firms in Germany. The fourth chapter analyzes the impact of the German electricity tax on the competitiveness of manufacturing firms. The fifth chapter analyzes the non-linear relationship between the EUA price and its fundamentals. The following sections briefly outline each of the four chapters.

1.6.1 Emissions trading and productivity: firm-level evidence from German manufacturing

Author: Benjamin Johannes Lutz

Other versions: Published as ZEW Discussion Paper No. 16-067

In the second chapter, I study the causal effect of the EU ETS on the productivity of German manufacturing firms. The existing empirical literature employing quasi-experimental approaches to investigate the impact of the EU ETS has focussed on the effects on emissions and other observable firm characteristics. So far, there is no study investigating the causal impact of the EU ETS - or any other emissions trading system - on a comprehensive and robust measure of economic performance.[1]

Addressing this gap in the literature, I estimate robust production functions for narrowly defined industries using administrative firm-level data from the German manufacturing sector. The underlying econometric model is based on the work by Ackerberg, Caves, and Frazer (2015) and De Loecker (2013). It allows for an endogenous dynamic productivity process and corrects for simultaneous changes in input use or productivity after a firm is regulated by the EU ETS.

After estimating the firm specific total factor productivity, I use different strategies to identify the causal effect of the EU ETS on the estimated productivity. Following the existing ex-post evaluation literature that is concerned with the EU ETS (Petrick and Wagner, 2014; Wagner, Muûls, Martin, and Colmer, 2014; Klemetsen, Rosendahl, and Jakobsen, 2016; Jaraitė and Di Maria, 2016; Calel and Dechezleprêtre, 2016), I exploit treatment variation created by the inclusion criteria of the EU ETS in order to isolate the causal effect of the EU ETS. In particular, I employ a difference-in-differences framework in order to identify and quantify the average treatment effect of the EU ETS on the productivity of regulated firms.

I exploit annual data from the German production census (Amtliche Firmendaten für Deutschland, AFiD) and the Cost Structure Survey for the period from 1999 to 2012

[1]See Section 1.4 for a detailed description of the literature.

gathered by the German statistical offices. The comprehensive firm-level information on inputs and output is merged with data from the EU Transaction Log, the official register of the EU ETS.

The results of the parametric difference-in-differences models suggest no significant negative effect of the EU ETS on productivity. In contrast, the EU ETS had a positive effect on productivity during the first compliance period. An alternative identification strategy based on a combination of the difference-in-differences framework and nonparametric nearest neighbor matching suggests that the EU ETS also had a positive effect on productivity during the second compliance period. A subsample analysis provides evidence that the effect of the EU ETS is heterogeneous across industries.

1.6.2 The impact of the EU ETS on economic performance of German manufacturing firms

Authors: Andreas Löschel, Benjamin Johannes Lutz, and Shunsuke Managi

In the third chapter, we investigate the effect of the EU ETS on the technical efficiency of manufacturing firms in Germany. First, we estimate a measure for economic performance that relates input use and produced output. In a subsequent step, we analyze the causal effect of the EU ETS on the economic performance.

The basic idea of this chapter is similar to how I proceed in the second chapter. The innovation lies in the alternative measure of economic performance. We depart from the concept of the mean production function. Instead, we estimate a stochastic production frontier for each two-digit industry following Aigner, Lovell, and Schmidt (1977). The frontier corresponds to the boundary of the production set of the industry that is determined by the most efficient firms. Based on the stochastic frontier, we compute firm-level technical efficiencies, i.e. the individual distance to the most efficient firms. The stochastic frontier approach relies on less stringent structural assumptions in comparison to the approach chosen in the second chapter. Instead, it is assumed that technical efficiency follows a certain type of distribution, in our case a half-normal distribution. After we obtain the firm-level technical efficiency, we employ a set of difference-in-differences models and nearest neighbor matching in order to identify the effect of the EU ETS on the regulated firms.

The data used in this chapter is different from the data used in the second chapter. The data requirements of the stochastic frontier model allow us to exclusively rely on the German production census (AFiD). In contrast to the Cost Structure Survey, AFiD comprises the universe of German manufacturing firms with more than twenty employees and thus provides more comprehensive picture of the German manufacturing sector.

None of our identification strategies provide evidence of a statistically significant negative effect of emissions trading on economic performance. In contrast, the results of the nearest neighbor matching suggest that the EU ETS had a positive impact on the economic performance of the regulated firms, especially during the first compliance period. A subsample analysis indicates that EU ETS increased the performance of treated firms in some two-digit industries while others remain unaffected.

1.6.3 The effect of electricity taxation on German manufacturing: a regression discontinuity approach

Authors: Florens Flues and Benjamin Johannes Lutz

Other versions: Published as ZEW Discussion Paper No. 15-013 and as OECD Environment Working Paper No. 88

Germany has taxed electricity use since 1999. In this chapter, we exploit discontinuities in the marginal tax rate in order to investigate the impact of electricity taxation on the competitiveness of manufacturing firms.

The German government granted reduced rates to energy intensive firms in the industrial sector to address potentially adverse effects on firms' competitiveness. Firms that use more electricity than certain thresholds established by legislation, pay reduced marginal tax rates. As a consequence, the marginal tax rate is a deterministic and discontinuous function of electricity use.

We identify and estimate the causal effects of these reduced marginal tax rates on the economic performance of firms using a sharp regression discontinuity design (Lee and Lemieux, 2010). In particular, we investigate how firms' turnover, exports, value added, investment, and employment responded to different marginal tax rates.

So far there is only one other study investigating the effect of the electricity price on German manufacturing. Gerster (2015) investigates the effect of a levy that is used to finance the feed-in tariff system for renewable energy. The levy is not applied to electricity intensive plants creating a discontinuity in the electricity price. Gerster (2015) employs a regression discontinuity design to identify the effect of the electricity price on electricity use and economic performance. He shows that the electricity use reacts to prices, however there is no evidence of a negative effect on economic performance. The analysis of Gerster (2015) is complementary to our work. The threshold of the discontinuity that he investigates applies for very large industrial electricity users, while the discontinuities created by the electricity tax are only relevant for small and medium-sized industrial electricity users.

Our econometric analysis relies on official micro-data at the plant and firm level (AFiD) collected by the German Federal Statistical Office that cover the entire manufacturing sector. We do not find any systematic, statistically significant effects of the electricity tax on firms' turnover, exports, value added, investment, or employment. We conduct several robustness checks to relax underlying assumptions and implement an alternative identification strategy. The results, however, remain unchanged. The findings suggest that gradually shifting the thresholds from which reduced tax rates apply may increase revenues for the government without adversely affecting the economic performance of firms.

1.6.4 Nonlinearity in cap-and-trade systems: the EUA price and its fundamentals

Authors: Benjamin Johannes Lutz, Uta Pigorsch, and Waldemar Rotfuß

Other versions: Published as ZEW Discussion Paper No. 13-001 and in Energy Economics Vol. 40 pp. 222-232

In the fourth chapter, we examine the nonlinear relation between the EUA price and its fundamentals, such as energy prices, macroeconomic risk factors and weather conditions.

When the underlying research strategy of this chapter was developed, the existing literature pointed to two facts. First, the EUA price time series is influenced by fundamentals. Second, the EUA price formation process is characterized by structural breaks that translate into this relationship. The paper was among the first to employ a nonlinear time series model that endogenously incorporates this changing nature of the EUA price formation process.

By estimating a Markov regime-switching model, we find that the relation between the EUA price and its fundamentals varies over time. In particular, we are able to identify a low and a high volatility regime, both showing a strong impact of the fundamentals on the EUA price. The most important EUA price drivers are changes on the stock market and energy prices. The gas price and a broad European equity index affect the EUA price positively in both regimes, while the coal price and the oil price have a significant, but also positive impact only during the high and the low volatility regime, respectively.

The high volatility regime is predominant in phases when economic activities are on a decrease or when institutional changes harm the confidence in the stringency of the EU ETS. This holds during the recession of 2008 and 2009, as well as during 2011 and 2012 when the debt crisis impaired the European economic outlook.

1.7 Avenues for future research

The German production census for manufacturing (AFiD) and the Cost Structure Survey offer a tremendous potential for empirical research - in particular with regard to topics related to energy and the environment. Apart from its manifold data on general firm and plant characteristics, AFiD contains very detailed information on energy and electricity use. This section of AFiD is far more detailed than production census data from other countries. It sheds light on topics such as fuel use, employment of renewables, and the generation of electricity through different kinds of technologies. Our work with this rich data set points to two directions for future research, that are briefly described in the following.

First, one of the advantages of AFiD is that it can be merged with external data at the plant and firm level. AFiD contains several unique identifiers, such as the commercial register number and the VAT number that can be used to match external firm-level data. Subsequently, data on the plant location can be used to match external plant-level data. The former procedure has been applied to create the combined data set that is the foundation for the second and the third chapter of this thesis. I merged AFiD with the EU Transaction Log, the official register of the EU ETS in order to obtain information on which firms are regulated. Gerster (2015) merges AFiD with data on the German feed-in-tariff levy. There are plenty of other external data sources that provide information on energy prices or environmental regulation faced by manufacturing firms. In combination with AFiD, these could be used to improve the understanding on how market-based instruments affect manufacturing firms. For instance, the European Pollutant Release and Transfer Register might be matched with AFiD in order to study its effects on the included firms.

Second, structural econometric models have scarcely been used to investigate the effects of energy prices and environmental regulation on manufacturing firms. The second chapter of this thesis is among the first using AFiD in order to estimate a structural econometric model. Based on the combination of a structural production function estimation and a set of difference-in-differences models, I find that the EU ETS had a positive effect on a revenue-based productivity measure. It is likely that firms increased product prices as reaction to the regulation by the EU ETS. The cost pass-through that might be reflected by the revenue-based productivity measure could be examined for at least some industries using a structural econometric model along similar lines to De Loecker, Goldberg, Khandelwal, and Pavcnik (2016). This is only one example of how the data set could be used for the estimation of structural econometric models and how they can contribute to the understanding of market-based instruments.

Chapter 2

Emissions trading and productivity: firm-level evidence from German manufacturing

2.1 Introduction

The increased likelihood of disasters linked to climate change, such as floodings, droughts, and forest fires, has put greenhouse gas reduction targets on policy agendas worldwide. The European Union (EU) acknowledged the importance of this endeavor by committing itself to cut greenhouse gas emissions by 20 percent until 2020 compared to 1990 figures. In 2005, the EU introduced the EU Emissions Trading System (EU ETS), a multinational cap-and-trade system, in order to regulate emissions from stationary industrial installations. The need for effective and efficient policy instruments to mitigate greenhouse gas emissions requires a thorough ex-post evaluation of existing policies. The analysis of the EU ETS and its impact on regulated firms offers the opportunity to better understand the mechanisms and consequences of market-based policy instruments.

First studies investigating the causal effects of the EU ETS show that it significantly reduced carbon dioxide emissions of regulated firms in many countries and spurred the development of new low carbon technologies throughout Europe (Petrick and Wagner, 2014; Wagner, Muûls, Martin, and Colmer, 2014; Klemetsen, Rosendahl, and Jakobsen, 2016; Calel and Dechezleprêtre, 2016). Despite its positive short and long term effects, the EU ETS has been the subject of significant political debate due to its potential adverse impact on the economic performance and competitiveness of regulated firms that compete on global markets.

From a theoretical point of view, it is not clear whether the EU ETS has an overall negative effect on the economic performance of regulated firms. The cost of complying with the EU ETS might decrease revenues and profits. However, according to the induced innovation hypothesis postulated by Hicks (1932), a relative change in input prices might create incentives to invest in new technology in order to reduce the use of the increasingly expensive input. In which way these adjustments affect economic performance depends on the production technology, input and output markets, and other factors that differ across firms. Porter (1991) and Porter and van der Linde (1995) put forward a stronger hypothesis specifically addressing the effects of environmental policy. They argue that properly designed regulation might not only increase the incentives to develop and adopt environmentally friendly technology but also consequently might affect competitiveness in a favorable manner.

Calel and Dechezleprêtre (2016) empirically investigate the induced innovation hypothesis and show that the EU ETS has a positive causal impact on low carbon innovation measured by patenting activities. These inventions, but also operational changes or investments in existing more efficient technologies might increase economic performance and create an advantage for regulated firms over competitors.

The aim of this study is to analyze the causal effect of the EU ETS on the economic performance of regulated firms. First, I estimate a structural production function model in order to obtain the total factor productivity as a robust and comprehensive measure of economic performance. Second, I isolate and quantify the effect of the EU ETS on the productivity of regulated firms by exploiting treatment variation that results from the design of the scheme.

The dispersion in productivity levels across firms and plants has been under investigation in many areas of economic research, such as industrial organization, labor, and trade (Syverson, 2011). Internal factors, such as input quality (Van Biesebroeck, 2003), managerial practice (Bloom and Van Reenen, 2007), or R&D activity (Doraszelski and Jaumandreu (2013)), but also external factors, such as trade competition (De Loecker, 2007) or market regulation (Knittel, 2002) drive firm-level productivity.[1]

Furthermore, environmental regulation has been identified as a driver of firm-level productivity. Greenstone, List, and Syverson (2012) investigate the impact of air quality

[1]See Syverson (2011) for a comprehensive survey on the productivity dispersion literature.

regulation on the productivity of U.S. plants. Regulations governing ozone, particulates, and sulfur dioxide decreased productivity, while regulations limiting carbon monoxide increased productivity. On average, the productivity of polluting plants decreased by 4.8 percent in regulated areas. Commins, Lyons, Schiffbauer, and Tol (2011) investigate the correlation between energy taxation and productivity during the period from 1997 to 2007. They use balance sheet data of European firms from the sectors manufacturing, energy, and transport. Exploiting industry-level variation in energy taxation, they find a positive correlation between taxation and firm-level productivity. In addition, they report a negative correlation between EU ETS participation and productivity for the first compliance period.

The origins of production function estimation date far back in economic literature. Marschak and Andrews (1944) investigate the challenges of estimating production functions using OLS. The observed input factors that enter the production function are chosen by firms. Therefore, unobserved firm specific determinants of production that are correlated with the choice of inputs are likely to bias OLS estimates. Recent advances have been made by Olley and Pakes (1996) (henceforth OP) who develop a structural econometric model of production that corrects for the described simultaneity bias using investments to proxy unobserved firm specific productivity shocks. Levinsohn and Petrin (2003) (henceforth LP) show that also static inputs can be used as proxies in the framework of the control function approach. Ackerberg, Caves, and Frazer (2015) (henceforth ACF) build on the basic idea to use observed firm characteristics to proxy the unobserved productivity, but suggest a more general and more robust estimation procedure. De Loecker (2013) enhances the ACF model by allowing for an endogenous productivity process.

The empirical strategy of this paper consists of two steps. First, I follow ACF and De Loecker (2013) and estimate a robust production function that allows the EU ETS to simultaneously influence input choice and productivity process in order to obtain an estimate for firm-level productivity. Secondly, I employ an identification strategy that relies on treatment variation caused by the design of the EU ETS in order to estimate the average treatment effect of the EU ETS on the productivity of regulated firms. In order to balance administrative cost, the EU exempts small and medium sized emitters from regulation by the EU ETS. As a consequence, there are regulated and unregulated firms within narrowly defined industries. I exploit this variation by estimating a variety of difference-in-differences models that isolate the effect of the EU ETS from confounding factors.

The empirical analysis is based on administrative firm data from Germany. I observe detailed annual firm-level information for the period from 1999 to 2012. The core of the dataset is the Cost Structure Survey (henceforth CSS) carried out by the German statistical offices. The CSS contains comprehensive annual information on output produced and inputs used by firms that operate in the manufacturing sector. It is the foundation of many governmental statistics and reports on the activities of the manufacturing sector.

The average treatment effect of the EU ETS on the productivity of regulated firms is statistically significant and ranges between 0.5 and 0.7 percent during the first compliance period depending on the set of control variables. The estimated treatment effect for the second compliance period is negative, but economically and statistically insignificant. Moreover, annual average treatment effects support the finding that the EU ETS had an impact during the first compliance period, while the estimated effects are not significantly different from zero for years of the second compliance period. In order to investigate the heterogeneity of the treatment effect across industries, I additionally estimate the difference-in-differences model for energy intensive two-digit industries with a sufficient number of regulated firms. The effect of the EU ETS on the productivity of firms from the food, paper, and chemical industry are not significantly different from zero. When estimating the model based on data for the industry producing basic metals, I find a significant positive effect for the first compliance period ranging between 2.4 and 2.9 percent.

In order to relax the parametric assumption of the difference-in-differences model, I follow Fowlie, Holland, and Mansur (2012) and combine the difference-in-differences framework with nonparametric nearest neighbor matching. The estimated treatment effects are statistically significant and positive for the first compliance period. In contrast to the parametric approach, the difference-in-differences matching model also provides evidence for a positive effect of the EU ETS during the second compliance period. The estimated treatment effects are slightly higher in comparison to the parametric approach and range between 1.5 and 2.7 percent for the first compliance period and 1.2 and 1.4 percent for the second compliance period, respectively. Annual treatment effect estimates based on the difference-in-differences matching model support the result that the EU ETS had a stronger effect during the first compliance period. For all models and specifications, I investigate pretreatment years and show that the productivity evolves in a parallel fashion across groups.

The structure of this paper is as follows: Section 2.2 describes the regulatory framework of the EU ETS and gives an overview of the empirical literature concerned with the causal effects of the EU ETS on firms' production and investment decisions. Section 2.3 describes the empirical production function estimation and the difference-in-differences model employed to identify the causal effect of the EU ETS. Section 2.4 describes the underlying official firm data and some descriptive statistics. Section 2.5 contains the parameter estimates of the empirical production functions and reports the results of the difference-in-differences model. Section 2.6 provides robustness checks with regard to heterogeneous treatment effects and functional assumptions. Section 2.7 discusses the results and concludes.

2.2 The EU ETS as a natural experiment

In order to achieve its emission targets in the framework of the Kyoto Protocol, the EU decided in 2003 to regulate greenhouse gas emissions from industrial installations by building an EU wide emissions trading system (European Parliament and Council, 2003). The resulting EU ETS was finally introduced in 2005 and currently regulates about 45 percent of the EU's greenhouse gas emissions caused by more than 11,000 installations in 31 countries.[2]

The EU ETS covers emissions from combustion installations and installations that run energy intensive production processes, such as oil refining, the production of metals, cement, lime, ceramics, bricks, glass, or paper. The design of the EU ETS excludes small and medium sized installations. For each of the listed processes, the European Commission (EC) defined specific capacity thresholds that determine the inclusion into the EU ETS.[3] The regulated installations have to undergo a continuous monitoring, reporting, and verification process. Once a year, they surrender allowances equivalent to their verified emissions.

[2] The EU ETS operates in the 28 member states of the EU as well as in Iceland, Liechtenstein, and Norway.

[3] See European Parliament and Council (2003) Annex I for details on the inclusion criteria.

The first compliance period of the EU ETS lasted from 2005 to 2007. It served as a pilot phase and was completely decoupled from the following compliance periods. As a consequence, allowances from the pilot phase were not eligible for surrender in later years. In 2005, the allowance price ranged between 20 and 30 euros. The price dropped and finally approached zero in 2007, when market participants realized that there was a massive oversupply on the market. The pilot phase was followed by second compliance period that coincided with the first commitment period of the Kyoto Protocol from 2008 to 2012. After the start of the second compliance period, the prices ranged between 20 and 30 euros, however a massive oversupply of allowances resulted in another price drop during the second half of 2008. As in previous years, the allowances were mostly allocated for free based on historic emissions. The plummeting demand due to the economic crisis and the massive use of project-based emission credits from the flexible mechanisms of the Kyoto Protocol did not meet the rather inelastic supply of allowances. The allowance price remained at a level of around 15 euros until a further shift in the second half of 2011 when the allowance price dropped below 10 euros. In contrast to 2007, the price did not converge toward zero since the allowances could also be used for compliance in subsequent years.[4]

Regulated firms can comply with the EU ETS in different ways. First, a regulated firm can surrender allowances to legitimate its emissions. This strategy is dominant as long as the allowance price is lower than the cost of abatement. Regardless of how the firm obtains the required allowances, the surrender of allowances negatively affects its economic performance. Alternatively, a regulated firm can abate greenhouse gas emissions through a change in input choice, an adjustment of the production technology, or the development of less emission intensive products. The effect of these abatement options on the firm's economic performance is less clear. A change in input choice, such as a fuel switch, comes at relatively low cost and might not have any additional effects on the firm. An adjustment of the production technology for example through investments in energy efficiency however might not only reduce the use of fuel but could also increase the overall efficiency of the firm. The development of new products also requires investments, but might create an advantage over competitors. From a theoretical perspective, it is not clear whether the EU ETS had a significant negative effect on firm performance or whether secondary effects

[4] See Figure 2.A.1 in Appendix 2.A for details.

of abatement measures increased economic performance.

A few very recent studies aim to contribute empirical evidence to the academic and public debates by investigating the causal effects of the EU ETS on the emissions and the economic performance of regulated firms. These studies exploit treatment variation that results from the inclusion criteria of the EU ETS. Since only large emitters are regulated, there are regulated and unregulated firms within narrowly defined industries that can be compared. The empirical evidence on the impact of the EU ETS on firm-level emissions and emission intensity is mixed. Studies using data from Germany, France, and Norway suggest that the EU ETS significantly reduced greenhouse gas emission (Petrick and Wagner, 2014; Wagner, Muûls, Martin, and Colmer, 2014; Klemetsen, Rosendahl, and Jakobsen (2016)).[5] Jaraitė and Di Maria (2016) do not find that the EU ETS significantly decreased firm-level emissions in Lithuania, but they find a significant negative effect on emission intensity. So far, there is no evidence that the EU ETS had a significant negative effect on indicators of economic performance. In contrast, Petrick and Wagner (2014) find a positive effect of the EU ETS on the revenues of regulated firms in Germany. Klemetsen, Rosendahl, and Jakobsen, 2016 find a positive effect on value added and labor productivity of regulated firms in Norway. Calel and Dechezleprêtre (2016) investigate the effect of the EU ETS on patenting. Their findings support that the EU ETS increased the number of low carbon patents developed by regulated firms. Bushnell, Chong, and Mansur (2013) examine how the EU ETS affects daily stock returns of European firms. They show that low allowance prices are associated with low stock prices for firms in both carbon and electricity intensive industries. Their results indicate that regulated firms profit from free allocation of allowances and a potential pass-through of environmental cost. Fabra and Reguant (2014) employ reduced-form and structural estimations in order to measure the pass-through of costs related to the EU ETS. Using Spanish electricity market data, they show that the environmental costs are almost completely passed through to electricity prices.

The aim of this study is to further investigate the impact of the EU ETS on the economic performance of regulated firms. Debates on potential detrimental effects of emissions trading on economic performance have accompanied the first two compliance

[5] Most of the studies investigating the causal effects of the EU ETS are still work under progress and thus not yet peer reviewed.

periods of the EU ETS. Large lobby groups such as the Federation of German Industries frequently point out that the EU ETS imposes high costs on regulated firms and thus threatens the competitiveness of European industries (The Federation of German Industries, 2016).[6] This paper gives nuance to these debates with sound empirical evidence on the causal effects of the EU ETS.

The existing ex-post evaluations investigate the effect of the EU ETS on revenues, value added, and labor productivity. These indicators do not provide a comprehensive picture of the firm specific economic performance. I employ a structural production function model in order to obtain a robust estimate of the firm-level total factor productivity. This measure for economic performance has been prominently used in the economic literature to investigate the origins of productivity dispersion.

Following the literature on the ex-post evaluation of the EU ETS, I exploit the inclusion criteria of the EU ETS to develop a sound identification strategy. I compare the firm specific productivity of firms that are regulated, i.e. belong to the treatment group, with the productivity of firms that are unregulated, i.e. belong to the control group, before and after the implementation of the EU ETS. I estimate the treatment effect employing parametric and nonparametric difference-in-differences models.

[6] The Federation of German Industries is an umbrella association representing the interests of more than 100,000 firms with about eight million employees.

2.3 Empirical strategy

The empirical analysis consists of two subsequent steps.[7] First, I estimate a robust production function that allows for an endogenous productivity process. It corrects for simultaneous changes in productivity and input use after a firm is regulated by the EU ETS. The resulting production function is used to recover a firm specific measure of productivity. Secondly, I assess the causal effect of the EU ETS on firm-level productivity by employing the difference-in-differences framework.

2.3.1 Production function estimation

Since the seminal paper of Marschak and Andrews (1944), it has been known that the estimation of a production function using OLS most likely leads to biased coefficients. The observed input factors that enter the production function - here, capital and labor - are chosen by the firm. If there is a firm specific determinant of production, that influences the input choice and is only visible to the firm itself, the OLS estimates will be biased.

OP develop a structural econometric model of production that corrects for the described simultaneity bias by using investments to proxy unobserved firm specific productivity shocks. LP enhance this approach and show that also static inputs, such as materials or energy use, can be used as proxy variables to control for productivity.[8] ACF build on the basic idea that underlies OP and LP, namely using observed firm characteristics to proxy the unobserved productivity, but suggest a more general and more robust estimation procedure. I estimate the production function following the approach proposed by ACF.

I assume a Cobb-Douglas production function that takes the following empirical log linear form for firm i at time t:

$$y_{it} = \beta_k k_{it} + \beta_l l_{it} + \omega_{it} + \varepsilon_{it}, \tag{2.1}$$

where y_{it} is the logarithm of gross value added, k_{it} and l_{it} denote the observable inputs capital and labor expressed in logs, and ω_{it} is the unobservable Hicks neutral productivity

[7] I follow the procedure applied by Braguinsky, Ohyama, Okazaki, and Syverson (2015) and employ a two-step approach in order to investigate the effect of the EU ETS on firm specific productivity changes. This procedure enables me to control for a large set of confounding factor in a difference-in-differences framework.

[8] Ackerberg, Benkard, Berry, and Pakes (2007) offer a more detailed technical description of the differences between these approaches.

term. The error term ε_{it} accounts for random shocks and measurement error and is assumed to be identically and independently distributed.

I assume, that the capital stock k at time t is determined by the investment i and the capital stock depreciation in $t-1$. The labor market regulation is quite strict in Germany by granting employees a period of notice that can last several months depending on employment relationship and industry. Therefore, I also treat labor as a dynamic input and assume that it is chosen between $t-1$ and t. OP, LP, and ACF assume the dynamics of productivity to evolve according to an exogenous first order Markov process. I follow De Loecker (2013) and consider a more general model, where I allow determinants of production to influence future productivity.[9] Accordingly, the evolution of the productivity dynamics is described by

$$\omega_{it+1} = f(\omega_{it}, \mathbf{z}_{it}) + \xi_{it}, \tag{2.2}$$

where the vector $\mathbf{z}$ collects determinants of production. The EU ETS might influence the production and investment decisions of a regulated firm and thus affect future productivity. Therefore, the effect of the EU ETS on productivity dynamics should be taken into account in order to prevent potential bias. I include a dummy variable indicating if a firm is regulated at time t. Furthermore, I add dummies for exports and R&D, since it has already been shown by De Loecker (2013), Doraszelski and Jaumandreu (2013), and Aw, Roberts, and Xu (2011) that both factors are important drivers of productivity. Accordingly, the functional relationship between ω_{it} and the determinants of production is governed by

$$f(\omega_{it}, \mathbf{z}_{it}) = \sum_{j}^{3} \theta_j \omega_{it}^j + \gamma_1 ets_{it} + \gamma_2 rnd_{it} + \gamma_3 exp_{it}, \tag{2.3}$$

where ets_{it}, rnd_{it}, and exp_{it} denote dummy variables, that indicate if a firm is regulated

[9] De Loecker (2013) includes exports into the first order Markov process in order to investigate the learning by exporting hypothesis. The idea to allow for endogenous productivity has been implemented in several structural econometric production models. Criscuolo and Martin (2009) examine the productivity of foreign owned plants in the United Kingdom, Doraszelski and Jaumandreu (2013) study the impact of R&D on productivity and Aw, Roberts, and Xu (2011) develop a structural model to shed light onto the joint impact of investments in R&D and exporting on productivity dynamics. Collard-Wexler and De Loecker (2015) measure the impact of technology choice on industry wide productivity in the U.S. steel industry. Braguinski, Ohyama, Okazaki, and Syverson (2015) investigate the effect of merger on productivity. All these studies feature the inclusion of additional production determinants into the first order Markov process that governs the productivity dynamics.

by the EU ETS, invests in R&D, or exports, respectively.

In order to deal with correlation between the observed inputs and productivity, I follow LP and rely on a firm's use of intermediate inputs m to control for unobserved productivity shocks that are captured by ω_{it}. I assume, that a firm's demand for the intermediate input is given by

$$m_{it} = g_t(k_{it}, l_{it}, \omega_{it}, \mathbf{z}_{it}). \tag{2.4}$$

I assume monotonicity of intermediate inputs in productivity and thus invert $g_t(\cdot)$ to obtain $\omega_{it} = h_t(k_{it}, l_{it}, m_{it}, \mathbf{z}_{it})$, the proxy for productivity in the empirical production function, i.e.

$$y_{it} = \beta_k k_{it} + \beta_l l_{it} + h_t(k_{it}, l_{it}, m_{it}, \mathbf{z}_{it}) + \varepsilon_{it}. \tag{2.5}$$

The parameters β_k and β_l are not identified in this equation, since they are both included in $h_t(\cdot)$. Following ACF, I nonparametrically estimate this equation to obtain the expected output $E[y_{it}|k_{it}, l_{it}, m_{it}, \mathbf{z}_{it}]$. As shown in Equation (2.6), the disposal of ε_{it} enables me to compute the productivity for any possible combination of β_k and β_l:

$$\omega_{it}(\beta_k, \beta_l) = E[y_{it}|k_{it}, l_{it}, m_{it}, \mathbf{z}_{it}] - \beta_k k_{it} - \beta_l l_{it}. \tag{2.6}$$

As in the ACF approach, I plug $\omega_{it}(\beta_k, \beta_l)$ into the law of motion of productivity in Equation (2.2) in order to obtain the error term given β_k and β_l, $\xi_{it}(\beta_k, \beta_l)$. The consequent moment conditions I employ to identify the parameters of the production function are described by

$$E\left[\xi_{it}(\beta_k, \beta_l) \middle| \begin{matrix} k_{it} \\ l_{it} \end{matrix}\right] = 0. \tag{2.7}$$

As explained above, I assume both capital and labor to be dynamic inputs and thus to be mean independent of ξ_{it}. I employ the generalized method of moments to estimate the parameters. Finally, I use the estimates for β_k and β_l to recover the implied productivity:

$$\hat{\omega_{it}} = y - \hat{\beta}_k k_{it} - \hat{\beta}_l l_{it}. \tag{2.8}$$

2.3.2 Disentangling the effect of the EU ETS

The second step of my empirical analysis is to identify and quantify the causal effect of the EU ETS by comparing changes in productivity across German manufacturing firms that are differentially affected by the EU ETS. Due to the inclusion criteria of the EU ETS, there are regulated and unregulated firms within narrowly defined industries allowing for a natural experiment framework. The baseline specification of the employed difference-in-differences model I estimate for the time period from 1999 to 2012 is

$$\begin{aligned} ln(Productivity_{it}) = \alpha_0 + \alpha_1 ETS_i + \alpha_2 ETS_i \times PhaseI_t \\ + \alpha_3 ETS_i \times PhaseII_t + \varphi_s + \delta_t + \lambda_{st} + u_{it}, \end{aligned} \tag{2.9}$$

where ETS_i indicates if a firm is regulated by the EU ETS, $PhaseI_t$ is equal to one for years during the first compliance period (2005-2007) and zero otherwise, and $PhaseII_t$ (2008-2012) is equal to one for years during the second compliance period (2008-2012). The inclusion of industry fixed effects φ_s adjusts for all constant unobserved determinants of productivity across industries. The year fixed effects δ_t control for superior trends in productivity in German manufacturing. λ_{st} denotes the full interaction terms between the industry and year fixed effects and nonparametrically absorbs within industry productivity trends. The error term u_{it} is assumed to be mean zero.

The parameters α_2 and α_3 on the interaction terms between ETS_i and the indicators for the two compliance periods $PhaseI_t$ and $PhaseII_t$ are the estimated effect of the EU ETS during Phase I and Phase II, respectively. In order to take into account observed and unobserved heterogeneity across regulated and unregulated firms, I enhance the baseline specification (Specification I) gradually. First, I add additional control variables, namely capital stock, employment, energy use, and material use (Specification II). Then, I add lagged indicators for export and R&D experience (Specification III) and firm fixed effects that adjust for all constant unobserved determinants of productivity across firms (Specification IV).

Key to the described identification strategy is the parallel trend assumption: I assume that in the absence of regulation by the EU ETS, trends in productivity evolve in a parallel fashion across groups conditional on the included control variables. In order to motivate this assumption, I investigate the development of productivity across treatment and control groups during pretreatment years. In particular, I estimate the following difference-in-differences model for the years before the announcement of the EU ETS, i.e. the period between 1999 and 2002:

$$\begin{aligned} ln(Productivity_{it}) = {} & \alpha_0 + \alpha_1 ETS_{it} + \alpha_2 ETS_{it} \times \mathrm{I}(t > 2000) \\ & + \varphi_s + \delta_t + \lambda_{st} + u_{it}, \end{aligned} \tag{2.10}$$

where $\mathrm{I}(t > 2000)$ equals one for the years 2001 and 2002. Therefore, the parameter α_2 is a placebo treatment effect. The parallel trend assumption would be violated, if α_2 is significantly different from zero. I apply this procedure to all aforementioned specifications of the difference-in-differences model.

The model described in Equation 2.9 provides the average treatment effects of the EU ETS on firm specific productivity for the first and the second compliance period. Within the two compliance periods, there was a high variation in the EUA price. High prices might have provoked stronger reactions by regulated firms and thus might have caused a heterogeneous treatment effect over time. In order to examine, if the effect of the EU ETS on productivity changes over time within the compliance periods, I estimate a modified difference-in-differences model that provides annual treatment effects:

$$\begin{aligned} ln(Productivity_{it}) = {} & \alpha_0 + \alpha_1 ETS_{it} + \sum_{k=2000}^{2012} \alpha_k ETS_{it} \times \mathrm{I}(t = k) \\ & + \varphi_s + \delta_t + \lambda_{st} + u_{it}, \end{aligned} \tag{2.11}$$

where I(t=k) is an indicator function associated with the year t. The estimated parameters for the years from 2000 to 2004 can be interpreted as placebo treatment effects. They will shed light onto the validity of the underlying assumptions of the difference-in-differences model. If these are significantly different from zero, the model fails to identify the effect of the EU ETS. The estimated parameters for the years from 2005 to 2012 will provide annual average treatment effects. I estimate this modified model applying the same specifications as described above (Specification I - IV).

2.4 Data and preliminary analysis

This study is based on official firm data from Germany. Combining different administrative data sources, I observe detailed annual firm-level information on general characteristics, cost structure, energy use, and EU ETS obligations for the time period from 1999 to 2012.

The core of the dataset is the Cost Structure Survey (CSS) carried out by the Federal Statistical Office and the Statistical Offices of the German Federal States. The CSS contains comprehensive annual information on output produced and inputs employed by firms that operate in the manufacturing sector. The CSS includes all German manufacturing firms with more than 500 employees. For firms with at least 20 and less than 500 employees, the statistical offices collect data from a large random sample. [10]The participation in the CSS is mandatory by law and results are checked for consistency and verified by the statistical offices. It is the foundation for many governmental statistics and reports on the activities of the manufacturing sector.

The random sample of firms with more than 20 and less than 500 employees is renewed once every few years - in the sample period at hand, the random sample has been drawn in 1999, 2003, 2008, and 2012. The random sample is stratified by the number of employees and industry affiliation. Firms with more than 20 and less than 500 employees are always surveyed if they belong to concentrated industries.

In order to add information on employment, exports, investments, and entry and exit to the CSS, I link the CSS with data from the German production census Official Firm Data for Germany (Amtliche Firmendaten für Deutschland - AFiD). The production census is also maintained by the German statistical offices and is obligatory for all manufacturing firms with more than 20 employees. Furthermore, I merge the CSS with the European Union Transaction Log (EUTL) in order to identify firms that are regulated by the EU ETS.[11] The resulting unbalanced panel comprises annual data of about 15,000 firms for the years from 1999 to 2012. Due to the dynamic structure of the model, I only consider firms that reported in at least two consecutive years. The industry classification system corresponds to ISIC Revision 4.

[10]Similar data sets from other countries have been used in the productivity dispersion literature. Doraszelski and Jaumandreu (2013) for example employ the Spanish equivalent in order to examine the effect of R&D on firm-level productivity

[11]See Appendix 2.A for more information on the merger of CSS, AFiD, and EUTL.

The measure for output is the firm's gross value added which is obtained from the CSS and deflated using two-digit ISIC deflators.[12] The labor input is constructed by taking the annual average of the number of employees reported monthly in the production census. The annual average offers a more detailed view on employment in comparison to the number of employees collected at the reporting date of the CSS. I use detailed investment data contained in the production census in order to compute the capital stock for each firm based on the perpetual inventory method.[13] The material expenditures stem from the CSS and are deflated by type using deflators for the manufacturing sector.[14] In addition, the firms participating in the CSS are asked to report R&D expenditures. These comprise the cost of internal R&D activities, but also joint activities with external research centers and laboratories. I consider a firm to conduct R&D activities, if the total R&D expenditures are positive. The production census provides export revenues on the firm-level. Analogously I consider a firm to be exporting if it reports positive export revenues.

In Table 2.1, I report descriptive statistics to characterize the group of firms regulated by the EU ETS and the group of unregulated firms. Firms regulated by the EU ETS are on average larger in terms of output produced and inputs used. This is due to the design of the EU ETS that only regulates large emitters. There is a relatively high number of small and medium sized firms in the data set. Therefore, statistics on the control group and the entire data set show positively skewed distributions on the main characteristics.[15] On average, the firms of the manufacturing sector were expanding until the economic crisis led to decreasing demand and contractions in 2009. From 2010 to 2012, the firms rapidly recovered and continued to grow on average.

Table 2.2 shows the distribution of regulated and unregulated firms across two-digit industries in the merged data set. While combustion installations that generate heat or power can be found in most industries, the firms that operate process regulated installa-

[12]The data on price indices can be retrieved from the web portal of the Federal Statistical Office: https://www-genesis.destatis.de/genesis/online Producer Price Index 61241-0004.

[13]This procedure has been used by many other papers estimating production functions, as for instance Olley and Pakes (1996). See Appendix 2.B for details on the computation of the capital stock.

[14]See footnote 11.

[15]Table 2.9 in Appendix 2.A reports detailed descriptive statistics on the entire data set showing percentiles and higher moments of the distributions.

tions are mainly concentrated in the industries manufacturing food (10), beverages (11), paper (17), coke and refined petroleum products (19), chemicals (20), pharmaceutical products (21), rubber and plastic (22), other nonmetallic mineral products (23), and basic metals (24).

Table 2.1: Descriptive statistics

	EU ETS firms			Unregulated firms		
	Mean	SD	N	Mean	SD	N
2000						
Gross value added (EUR 1,000)	270,956	930,476	339	19,323	186,228	14,768
Output (EUR 1,000)	623,951	2,348,248	338	45,760	415,962	14,660
Capital stock (EUR 1,000)	277,470	851,531	339	15,825	117,142	14,700
Energy use (MWh)	1,139,299	4,880,839	339	21,816	293,735	14,767
Number of employees	2,339	8,429	339	248	1,661	14,767
R&D expenditure (EUR 1,000)	35,334	215,372	339	1,636	41,329	14,768
Exports (EUR 1,000)	417,213	2,165,096	339	21,343	324,984	14,767
2005						
Gross value added (EUR 1,000)	296,848	1,173,846	383	19,685	178,177	13,475
Output (EUR 1,000)	711,907	2,778,015	381	50,103	443,246	13,223
Capital stock (EUR 1,000)	278,734	930,450	383	16,143	117,655	13,389
Energy use (MWh)	1,291,501	4,469,597	383	19,937	169,754	13,286
Number of employees	2,260	8,013	383	241	1,622	13,474
R&D expenditure (EUR 1,000)	44,205	253,976	383	1,913	45,681	13,475
Exports (EUR 1,000)	503,908	2,511,781	383	26,790	370,656	13,474
2010						
Gross value added (EUR 1,000)	270,600	1,050,857	440	18,274	201,533	14,959
Output (EUR 1,000)	696,410	3,015,317	438	47,815	487,422	14,812
Capital stock (EUR 1,000)	254,298	814,679	440	14,310	109,372	14,884
Energy use (MWh)	1,595,464	6,518,306	440	17,217	131,277	14,818
Number of employees	1,936	7,032	440	220	1,460	14,944
R&D expenditure (EUR 1,000)	42,431	258,910	440	1,715	41,345	14,959
Exports (EUR 1,000)	540,497	3,056,996	440	26,037	372,078	14,945

Notes: Gross value added, output (production value), wages and salaries, R&D expenditure, exports and capital stock are denoted in EUR 1,000. Energy use is denoted in MWh. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - Cost Structure Survey, AFiD-Panel Industrial Units, and AFiD-Module Use of Energy, own calculations.

Table 2.2: Number of observations by industry - the Cost Structure Survey

		2005		2008		2012	
NACE	Industry	Total	Regulated	Total	Regulated	Total	Regulated
10	Food products	1,528	44	1,903	52	1,888	53
11	Beverages	285	8	260	10	245	11
12	Tobacco products	22	1	21	2	21	2
13	Textiles	424	5	414	7	333	4
14	Wearing apparel	268	0	215	0	185	0
15	Leather and related products	118	0	103	0	84	0
16	Wood and products of wood and cork	361	11	458	19	377	16
17	Paper and paper products	359	63	435	72	435	78
18	Printing and reproduction of recorded media	330	2	359	1	329	3
19	Coke and refined petroleum products	44	15	47	17	45	16
20	Chemicals and chemical products	729	49	854	56	846	54
21	Pharmaceutical products	184	8	201	8	185	6
22	Rubber and plastic products	806	10	978	9	812	14
23	Other nonmetallic mineral products	733	100	792	117	687	119
24	Basic metals	545	32	642	32	646	35
25	Fabricated metal products	1,522	2	2,149	4	1,907	2
26	Computer, electronic and optical products	694	5	831	4	694	4
27	Electrical equipment	808	4	970	4	980	4
28	Machinery and equipment n.e.c.	2,226	6	2,585	8	2,138	7
29	Motor vehicles, trailers, and semitrailers	663	10	701	9	541	9
30	Other transport equipment	257	5	238	5	212	5
31	Furniture	382	0	397	0	351	0
32	Other manufacturing	452	3	587	3	547	2
33	Repair and installation of mach. and equip.	118	0	108	0	623	1
–	Total	13,858	383	16,248	439	15,111	445

Notes: Number of firms for the first year of Phase I of the EU ETS (2005), the first year of Phase II (2008) and the last year of Phase II (2012) that is also the last year I observe. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - Cost Structure Survey and AFiD-Panel Industrial Units, own calculations.

2.5 Results

In this section, I first report the production function estimates based on the ACF approach I described in Section 2.3.1. Secondly, I will show the results of the difference-in-differences model I outlined in Section 2.3.2 in order to shed light onto the causal effect of the EU ETS on firm-level productivity.

2.5.1 Production function estimates

I estimate value added production functions for two-digit industries within manufacturing using data for the time period from 1999 to 2012. Table 2.3 reports the results of the ACF model along with estimates based on standard OLS regressions, the total number of firms and the total number of observations for the entire sample period.

The ACF model is implemented employing the general method of moments procedure shown in Section 2.3.1. The standard errors are clustered at the firm-level and obtained by applying the block bootstrap algorithm treating each set of firm observations together as an independent and identical draw from the population of firms.[16] The block bootstrap algorithm takes into account that the multiple observations of a firm are correlated over time in some unknown way and corrects for the two-step nature of the general method of moments estimator.

The number of firms within the industries ranges between 253 (leather and related products) and 6,760 (machinery and equipment) firms during the sample period from 1999 to 2012, while the total number of observations ranges between 1,582 (leather and related products) and 31,085 (machinery and equipment). The precision of the estimates tends to increase with the number of observations. All estimated coefficients of the ACF model are statistically significant at conventional levels.

The production functions vary significantly across industries reflecting the heterogeneity within the manufacturing sector. All industries have in common that the coefficient on labor is larger than the coefficient on capital. A comparison of the ACF and the OLS parameters shows that for most industries, the ACF coefficient on capital is larger, while

[16] This approach has been first proposed by LP in this context. Subsequent studies that apply LP or ACF follow this strategy, see e.g. De Loecker and Warzinski (2012), De Loecker (2013), and Collard-Wexler and De Loecker (2015). More information on the bootstrap can be found in Horowitz (2001).

Table 2.3: Output elasticities

				OLS estimates		ACF estimates	
NACE	Industry	# Firms	# Observ.	Capital	Labor	Capital	Labor
10	Food products	4,342	22,981	0.349 (0.004)	0.643 (0.005)	0.369 (0.024)	0.541 (0.015)
11	Beverages	664	3,889	0.137 (0.013)	0.985 (0.020)	0.181 (0.065)	0.949 (0.044)
13	Textile	1,089	6,030	0.183 (0.006)	0.885 (0.010)	0.195 (0.022)	0.867 (0.038)
15	Leather and related products	253	1,582	0.249 (0.012)	0.834 (0.019)	0.237 (0.035)	0.839 (0.055)
16	Wood and products of wood and cork	1,261	5,640	0.175 (0.006)	0.884 (0.010)	0.180 (0.028)	0.844 (0.038)
17	Paper and paper products	1,017	5,532	0.214 (0.007)	0.867 (0.010)	0.289 (0.034)	0.788 (0.052)
18	Printing and reproduction of recorded media	1,207	4,704	0.152 (0.007)	0.900 (0.010)	0.257 (0.087)	0.613 (0.196)
20	Chemicals and chemical products	1,707	10,311	0.241 (0.006)	0.820 (0.008)	0.258 (0.024)	0.779 (0.039)
22	Rubber and plastic products	2,652	11,864	0.183 (0.004)	0.881 (0.006)	0.198 (0.014)	0.865 (0.020)
23	Other nonmetallic mineral products	2,061	10,836	0.249 (0.005)	0.792 (0.007)	0.225 (0.016)	0.820 (0.019)
24	Basic metals	1,269	8,088	0.176 (0.006)	0.868 (0.009)	0.190 (0.026)	0.855 (0.036)
25	Fabricated metal products	5,791	24,835	0.148 (0.003)	0.926 (0.004)	0.144 (0.008)	0.925 (0.013)
27	Electrical equipment	2,610	11,807	0.161 (0.005)	0.909 (0.007)	0.209 (0.054)	0.834 (0.079)
28	Machinery and equipment n.e.c.	6,760	31,085	0.092 (0.003)	0.994 (0.004)	0.091 (0.007)	0.988 (0.010)
29	Motor vehicles, trailers, and semitrailers	1,553	8,451	0.164 (0.006)	0.897 (0.008)	0.164 (0.006)	0.886 (0.023)
31	Furniture	1,256	5,544	0.156 (0.006)	0.927 (0.009)	0.168 (0.050)	0.853 (0.089)

Notes: All parameter estimates are significant at the 5 percent level. Standard errors are computed by employing the block bootstrap algorithm with 500 replications. I employ cluster bootstrap to obtain standard errors. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - Cost Structure Survey and AFiD-Panel Industrial Units own calculations.

the ACF coefficient on labor is smaller. This is in line with the findings of OP and LP. They show that the endogeneity of the input choice results in an upward or downward bias depending on how fast a firm can adjust the input use. OLS coefficients on relatively inflexible inputs such as capital tend to be biased towards zero, while coefficients on relatively flexible inputs are positively biased. The returns to scale range from 0.87 (printing and reproduction of recorded media) to 1.13 (beverages).

I compute the firm-level productivity as the residual from the production function as described in Equation 2.8. Figure 2.1 shows the indexed mean productivity for two-digit industries within manufacturing. The mean productivity evolves quite differently over time across two-digit industries again reflecting the heterogeneity across industries within manufacturing. While some industries record increasing mean productivity (e.g. food products, electrical equipment, machinery and equipment, and motor vehicles, trailers, and semitrailers), others show decreasing mean productivity (e.g. printing and reproduction of recorded media and basic metals, and furniture).

The economic crisis did not affect manufacturing in Germany as seriously as in other European countries. However, demand for German goods decreased significantly in 2009. This development is also reflected in a drop in productivity of most two-digit industries. Firms cannot smoothly adapt their input choice to demand shocks. Therefore, a rapid decline in demand decreases capacity utilization and consequently productivity.

The different developments in productivity across industries are taken into account in the second stage of my empirical analysis. The various specifications of the difference-in-differences model always include industry and year fixed effects as well as a full set of interaction terms. In this way, the model nonparametrically captures heterogeneous developments in productivity across industries and over time.

Figure 2.1: Indexed mean productivity (base year 1999)

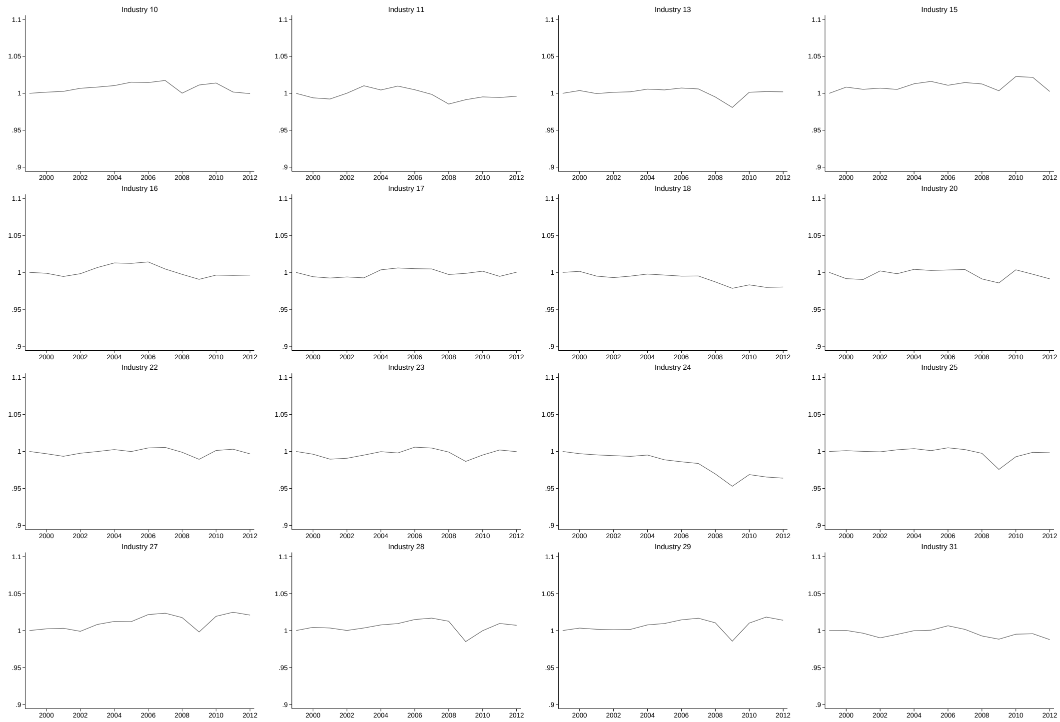

Notes: Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - Cost Structure Survey and AFiD-Panel Industrial Units, own calculations.

2.5.2 Estimated treatment effects

Table 2.4 reports the results of the difference-in-differences model described in Section 2.3.2. The first column shows the estimates of Specification I that includes fixed effects and full interaction terms on industry and year. Columns two and three report the results for Specification II and Specification III that include an enhanced set of control variables. The results of Specification IV that adds firm fixed effects are shown in the last column. The displayed treatment effects can be interpreted as semi-elasticities.

Bertrand, Duflo, and Mullainathan (2004) point out that conventional standard errors for difference-in-differences applications with long time series and a high serial correlation in the outcome variable are inconsistent. Since the considered time series is rather long (14 years) and productivity is highly persistent, I refrain from applying conventional standard errors. According to Bertrand, Duflo, and Mullainathan (2004) the bootstrap performs well if the cross section is sufficiently large and the serial correlation in the data is taken into account. I follow their recommendation and employ the block bootstrap algorithm in order to obtain adequate standard errors for the estimated treatment effects clustered at the firm-level.

Specifications I, II, and IV show a significant positive effect of the EU ETS on firm-level productivity ranging between 0.5 and 0.7 percent during the first compliance period. The estimated treatment effects for the second compliance periods are negative, but rather small and statistically insignificant. In order to examine the parallel trend assumption, I estimate the four specifications for the pretreatment time period from 1999 to 2002 treating 2001 as the implementation year of the EU ETS. The results of the pretreatment analysis are report in the third row of Table 2.4. None of the estimated placebo effects is statistically significant. Consequently, I fail to reject the hypothesis that the key identifying assumption holds strongly supporting the validity of the difference-in-differences approach in this setting.

In addition to the average treatment effect for the entire compliance period, I also estimate the annual effects of the EU ETS as shown in Equation 2.11 in order to investigate variations in the impact over time. Figure 2.2 displays the results of the annual treatment effects model. The horizontal line denotes the treatment effect while the horizontal bar denotes the twofold standard deviation. Per year, four bars are shown corresponding to the four specifications outlined above. Most estimated treatment effects are statistically insignificant. However, the annual treatment effects of Phase I are slightly higher in comparison to the pretreatment and Phase II years.

Table 2.4: Difference-in-differences treatment effects

	I	II	III	IV
Phase I	0.007**	0.005*	0.005	0.007**
	(0.003)	(0.003)	(0.003)	(0.003)
Phase II	-0.000	-0.002	-0.001	-0.001
	(0.003)	(0.003)	(0.003)	(0.003)
Pretreatment analysis	0.005	0.004	0.005	(0.004)
	(0.003)	(0.003)	(0.003)	(0.003
Fixed effects				
Industry	×	×	×	
Year	×	×	×	×
Industry year interaction terms	×	×	×	×
Firm				×
Additional controls				
Capital, labor, energy use, and materials		×	×	×
Indicator for export and RnD experience			×	×
# Firms (1999-2012)	34,373	34,215	32,302	32,302
# Observations (1999-2012)	173,178	172,065	153,823	153,823

Notes: Standard errors are computed by employing the block bootstrap algorithm with 500 replications. *** denotes significance at the 99 percent level. ** denotes significance at the 95 percent level. * denotes significance at the 90 percent level. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - Cost Structure Survey, AFiD-Panel Industrial Units, and AFiD-Module Use of Energy, own calculations.

Figure 2.2: Annual treatment effects

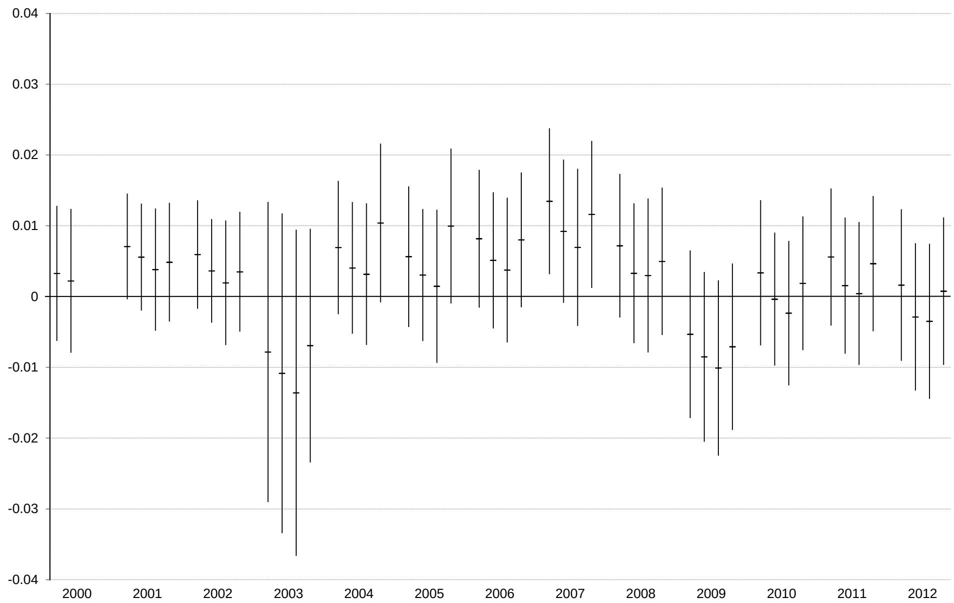

Notes: Annual treatment effects and confidence bands (2 times standard error) for Specification I - Specification IV. Specification III and Specification IV can only be estimated for the period from 2000 to 2012. I employ the block bootstrap algorithm with 500 replications to obtain robust standard errors. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - Cost Structure Survey and AFiD-Panel Industrial Units, own calculations.

2.6 Robustness checks

2.6.1 Heterogeneous treatment effects

The industries within the manufacturing sector differ with regard to many aspects. They produce very different goods, face different market conditions on input and output markets and face different kinds of regulation. As a consequence, the effect of the EU ETS on the regulated firms might vary across industries. The average treatment effect over all industries therefore does not provide the full picture of the impact of the EU ETS.

For this reason, I analyze the effect of the EU ETS at the two-digit industry level. I focus on the industries manufacturing food products (10), paper and paper products (17), chemicals and chemical products (20), rubber and plastic products (22), other nonmetallic mineral products (23) as well as basic metals (24). These industries provide a sufficient number of observations in the group of the regulated firms (see Table 2.2). I estimate the difference-in-differences model depicted in Equation 2.9 and further investigate the validity of the parallel trend assumption by estimating a placebo treatment effect during the pretreatment period from 1999 to 2002 as described by Equation 2.10. The standard errors are again obtained by applying the block bootstrap algorithm.

Table 2.5 reports the results for Specification I and II and Table 2.6 shows the results for Specification III and IV, respectively. Considering the first compliance period, the estimated treatment effects for the industries manufacturing food (10), paper (17), and chemicals (20) are mostly positive, but statistically insignificant. During the second compliance period, the effect of the EU ETS was statistically insignificant, but the signs varied across the three industries. The estimated treatment effect was positive for the industries producing food (10) and chemicals (20), while the effect was negative for the paper industry (17). The results for the industries producing rubber and plastic (22) and other nonmetallic mineral products (23) indicate that the parallel trend assumption is violated for some specifications. Therefore, I refrain from interpreting the results for these industries. The EU ETS had a significant positive effect on the firms of the industry producing basic metals (24) during the first compliance period. The effect ranges between 2.4 and 2.9 percent. During the second compliance period, the EU ETS did not significantly influence the productivity of the regulated firms. For this industry, the pretreatment analysis supports the parallel trend assumption.

The subsample analysis sheds light on the heterogeneity of the treatment effect, however, this empirical strategy comes along with a reduction of the sample size. The number of firms ranges from 1,707 (paper industry, 17) to 4,342 (food industry, 10) in the period from 1999 to 2012, while the total number of observation ranges from 10,310 (paper industry, 17) to 22,981 (food industry, 10). As a consequence, the precision of the estimates decreases in comparison to the analysis using the full sample.

The results of the subsample analysis support the hypothesis that the effect of the EU ETS is not homogeneous across industries.

Table 2.5: Difference-in-differences treatment effects - subsample analysis (I/II)

		I			II		
NACE	Industry	Pre-treatment	Phase I	Phase II	Pre-treatment	Phase I	Phase II
10	Food products	0.019 (0.017)	0.002 (0.009)	0.002 (0.010)	0.019 (0.018)	0.001 (0.009)	0.003 (0.010)
17	Paper and paper products	0.013 (0.010)	0.009 (0.012)	-0.004 (0.013)	0.014 (0.010)	0.008 (0.012)	-0.005 (0.013)
20	Chemicals and chemical products	-0.001 (0.005)	0.002 (0.007)	0.007 (0.007)	-0.001 (0.005)	0.002 (0.006)	0.005 (0.006)
22	Rubber and plastic products	0.002 (0.007)	0.005 (0.005)	0.011* (0.006)	0.014 (0.010)	0.001 (0.006)	0.007 (0.005)
23	Other non-metallic mineral products	-0.004 (0.004)	-0.002 (0.005)	-0.001 (0.004)	-0.006 (0.004)	-0.004 (0.004)	-0.003 (0.004)
24	Basic metals	0.001 (0.004)	0.029*** (0.007)	-0.005 (0.010)	-0.001 (0.004)	0.026*** (0.007)	-0.009 (0.010)

Notes: Standard errors are computed by employing the block bootstrap algorithm with 500 replications. *** denotes significance at the 99 percent level. ** denotes significance at the 95 percent level. * denotes significance at the 90 percent level. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - Cost Structure Survey, AFiD-Panel Industrial Units, and AFiD-Module Use of Energy, own calculations.

Table 2.6: Difference-in-differences treatment effects - subsample analysis (II/II)

		III			IV		
NACE	Industry	Pre-treatment	Phase I	Phase II	Pre-treatment	Phase I	Phase II
10	Food products	0.022 (0.018)	0.004 (0.010)	0.006 (0.011)	0.021 (0.017)	0.012 (0.009)	0.008 (0.010)
17	Paper and paper products	0.012 (0.009)	0.008 (0.013)	-0.006 (0.014)	0.013 (0.009)	0.012 (0.013)	-0.005 (0.007)
20	Chemicals and chemical products	0.002 (0.004)	-0.001 (0.006)	0.003 (0.006)	0.001 (0.004)	0.001 (0.006)	0.005 (0.007)
22	Rubber and plastic products	0.026** (0.012)	-0.002 (0.007)	0.004 (0.006)	0.020* (0.010)	0.002 (0.006)	0.008 (0.006)
23	Other non-metallic mineral products	-0.007 (0.004)	-0.004 (0.005)	-0.001 (0.004)	-0.007** (0.004)	0.002 (0.004)	-0.002 (0.004)
24	Basic metals	0.002 (0.004)	0.025*** (0.007)	-0.011 (0.010)	-0.002 (0.004)	0.024*** (0.008)	-0.013 (0.010)

Notes: Standard errors are computed by employing the block bootstrap algorithm with 500 replications. *** denotes significance at the 99 percent level. ** denotes significance at the 95 percent level. * denotes significance at the 90 percent level. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - Cost Structure Survey, AFiD-Panel Industrial Units, and AFiD-Module Use of Energy, own calculations.

2.6.2 Relaxing functional form assumptions

The parametric difference-in-differences model described in Section 2.3.2 relies on functional form assumptions on the treatment and outcome model that might affect the estimated treatment effect. When including observable firm characteristics into the difference-in-differences model, I implicitly assume a linear relationship between the control variables, the treatment, and the outcome. In order to relax this assumption, I follow Fowlie, Holland, and Mansur (2012) and combine the difference-in-differences framework with non-parametric nearest neighbor matching. Based on the firm characteristics I observe, I match treated firms with similar untreated firms. The resulting control group only contains firms that are similar to the firms of the treatment group.

Following Heckman, Ichimura, and Todd (1997) and Heckman, Ichimura, Smith, and Todd (1998) the difference-in-differences matching estimator is described by

$$\begin{aligned}\hat{\tau} = \frac{1}{N}\sum_{j\in I_1}\Big\{&(ln(Productivity)_{jt'}(1) - ln(Productivity)_{jt^0}(0)) \\ &- \sum_{k\in I_0} w_{jk}(ln(Productivity)_{kt'}(0) - ln(Productivity)_{kt^0}(0))\Big\}.\end{aligned} \tag{2.12}$$

The set of firms regulated by the EU ETS is defined as I_1, while the unregulated firms are collected in set I_0. There are N regulated firms indexed by j, the unregulated firms are indexed by k. The weight w_{jk} takes the value one, if a firm of the control has been matched and zero otherwise. Following Fowlie, Holland, and Mansur (2012), I identify the nearest neighbor using the Mahalanobis distance to measure similarity between firms. I perform matching with replacement linking each treated firm with one and five similar firms of the control group. Similar firms are identified using information of the pretreatment year 2000 on output, capital stock, labor, energy use, material use as well as indicators for export and R&D experience. Due to the strong heterogeneity across industries, I exactly match on the two-digit industry classification.

Similarly to the parametric difference-in-differences approach, I assume that the mean productivities of treatment and control group evolve in a parallel fashion over time in the absence of the regulation by the EU ETS. In order to investigate the validity of this assumption, I conduct a pretreatment analysis using data from 1999 to 2002 with a placebo treatment in 2001.

Table 2.7 reports the results of the pretreatment analysis and the estimated treatment effects based on the difference-in-differences matching estimator. The estimates are based on a comparison between the pretreatment period from 2001 to 2002 with each compliance period of the EU ETS. I compute standard errors that are robust with respect to autocorrelation and heteroscedasticity following Abadie and Imbens (2006). The placebo treatment effects estimated for the year 2001 are close to zero and statistically insignificant. The treatment effect estimated for the first compliance period is statistically different from zero. When matching with the nearest neighbor, I obtain an average treatment effect of 2.7 percent. The preferred specification of the parametric difference-in-differences model provides a productivity increasing effect of 0.7 percent. Adding the five closest neighbors to the control group reduces the treatment effect to 1.5 percent. Also for the second compliance period, the nearest neighbor matching shows significantly positive estimates. The treatment effect ranges between 1.2 percent (one neighbor) and 1.4 (five neighbors). These results differ from the results of my main model that does not provide any evidence for a significant effect of the EU ETS during the second compliance period. These differing outcomes might be explained by the different designs of the two identification strategies. Applying the matching algorithm, I avoid the functional assumptions of the parametric difference-in-differences model and I only compare the regulated firms with very similar unregulated firms. Furthermore, I am only able to compare firms that stay in the sample during the considered time periods from 1999 to 2008 and from 1999 to 2012, respectively.

Apart from the average treatment effect for each compliance period, I also estimate annual treatment effects in order to investigate the development of the treatment effect over time. The pretreatment year 2000 serves as the base year for this approach. For the pretreatment years from 2001 to 2004, I expect the estimated treatment effects not to be statistically different from zero. Figure 2.3 shows the annual treatment effects and the corresponding confidence bands. From 2001 to 2004, the estimated treatment effect is not significantly different from zero. During the first compliance period, the confidence bands slightly widen, however the estimated treatment effects are statistically different from zero. During the first compliance period, the estimated treatment effect ranges between 1.4 percent (2006) and 3.8 percent (2007) when matching with the nearest neighbor and between 1.1 percent (2006) and 1.8 percent (2008) when matching with the five nearest neighbors. The estimated treatment effects for the second compliance period are closer to zero and statistically insignificant.

The results of the difference-in-differences nearest neighbor matching approach support the results of the parametric difference-in-differences model.

Table 2.7: Nonparametric difference-in-differences treatment effects

	Pretreatment analysis	Phase I	Phase II
One neighbor	0.006	0.027***	0.012*
	(0.004)	(0.009)	(0.006)
Five neighbors	0.004	0.015***	0.014***
	(0.003)	(0.005)	(0.005)
# Observations	11,609	3,212	6,757

Notes: The computed standard errors are based on Abadie and Imbens (2006). *** denotes significance at the 99 percent level. ** denotes significance at the 95 percent level. * denotes significance at the 90 percent level. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - Cost Structure Survey, AFiD-Panel Industrial Units, and AFiD-Module Use of Energy, own calculations.

Figure 2.3: Annual treatment effects - NN matching

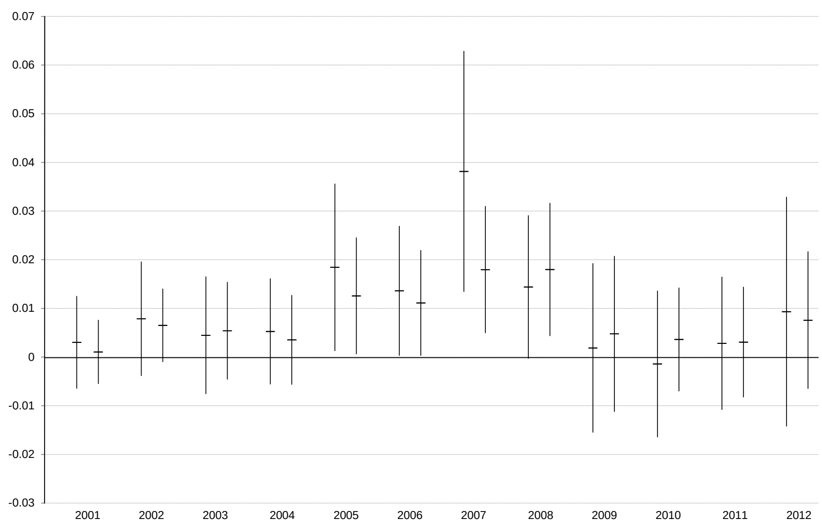

Notes: Annual treatment effects and confidence bands (2 times standard error) for nearest neighbor matching with one and five neighbors. As base year for the difference-in-differences approach serves the year 2000. The computed standard errors are based on Abadie and Imbens (2006). Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - Cost Structure Survey and AFiD-Panel Industrial Units, own calculations.

2.7 Concluding discussion

Debates on potential detrimental effects of emissions trading on economic performance have accompanied the first two compliance periods of the EU ETS. This paper investigated the causal effect of the EU ETS on the total factor productivity of regulated firms. The productivity is measured as the residual from a structural production function estimate that is robust with regard to endogeneity and allows the EU ETS to simultaneously influence input choice and the dynamic productivity process. I identify the effect of the EU ETS on firm-level productivity by exploiting treatment variation that occurs due to the design of the EU ETS. In order to release small and medium sized emitters of carbon dioxide from the regulatory burden, only large emitters have been included into the EU ETS. I examine changes in the firm-level productivity by comparing regulated and unregulated firms before and after the implementation of the EU ETS. In particular, I estimate a variety of parametric difference-in-differences models including different sets of explanatory variables and fixed effects in order to eliminate the influence of potential confounding factors.

The estimated treatment effects for the first compliance period of the EU ETS are mostly statistically significant and positive. No effect of the EU ETS could be observed for the second compliance period using the parametric difference-in-differences model. Estimated treatment effects based on a combination of the difference-in-differences framework and nearest neighbor matching support the findings for the first compliance period, but also show a significant positive effect of the EU ETS on productivity during the second compliance period.

So far, there is no scientific study investigating the effect of the EU ETS on the total factor productivity of regulated firms. The very recent ex-post evaluations of the EU ETS mostly find a significant reduction of carbon dioxide emissions (Petrick and Wagner, 2014; Wagner, Muûls, Martin, and Colmer, 2014; Klemetsen, Rosendahl, and Jakobsen, 2016) and emission intensity (Petrick and Wagner, 2014; Jaraitė and Di Maria, 2016) caused by the EU ETS. With regard to other firm characteristics, Petrick and Wagner, 2014 find a positive of the EU ETS on the revenue of German firms and Klemetsen, Rosendahl, and Jakobsen, 2016 find a positive effect on value added and labor productivity. Jaraitė and Di Maria (2016) provide empirical evidence for a positive effect of the EU ETS on the renewal of installed capital stock. Wagner, Muûls, Martin, and Colmer, 2014 find that the EU ETS had a negative effect on employment. Calel and Dechezleprêtre (2016) investigate the effect of the EU ETS on patenting and find that it spurred low-carbon patents.

The results of this study are basically in line with these findings that are based on data from different European countries. Investments in more efficient capital stock triggered by the EU ETS could have reduced the use of static inputs such as energy and labor and thus might have increased the overall productivity. Innovative new products that require less energy use in the production process might have increased productivity and output as well.

There are several factors that might have influenced the analysis and should be kept in mind when discussing the results and their implications. First, I would like to point to the fact that the productivity measure employed in this study is revenue based. I use specific two-digit price indicators to deflate outputs produced and inputs employed, however there might be still changes driven by price developments. This has implications for the interpretation of the causal effect of the EU ETS. If firms are able to pass additional cost on to customers, then the higher output price due to a cost pass-through could be reflected as a productivity gain. As shown by Bushnell, Chong, and Mansur (2012) and Fabra and Reguant (2014), emissions costs are passed through to electricity prices by utilities. It is likely that at least in some industries of the manufacturing sector firms are also able to pass through costs to product prices. A solution to this problem would be to estimate a quantity based productivity measure based on physical units of inputs and outputs. The estimation of a quantity based production function for multi-product firms is in general feasible (De Loecker, Goldberg, Khandelwal, and Pavcnik, 2016). However, this approach requires a sufficient number of single-product firms within narrowly defined industries. Using only German data, I do not observe enough single-product firms in order to estimate such a model.

A further issue that might influence the results of my analysis is the exit of firms caused by the EU ETS. If the EU ETS forced unproductive regulated firms out of the market, then the average productivity would increase in the group of the regulated firms even if the active firms remained on the same productivity level. Regrettably, the CSS is an unbalanced panel similar to the data sets used in many other studies of the productivity dispersion literature, such as Doraszelski and Jaumandreu (2013). The design of the

CSS impedes a clear distinction between the exit from the market and the end of the obligation to participate in the survey. However, examining the EUTL data I do not observe significant attrition among the EU ETS regulated firms.

The second compliance period of the EU ETS coincides with the world economic crisis. In 2009, output produced and inputs employed decreased in the German manufacturing sector. Firms recovered quickly in the subsequent years, however consequences of the crisis might conflate with the effect of the EU ETS during the second compliance period. I assume that the crisis did not affect firms from the treatment and control group differently conditional on fixed effects and firm characteristics. Among other control variables, I include capital stock, employment, energy use, and material use into the difference-in-differences model in order to capture the variation caused by the crisis.

In this study, I focus on the direct effect of the EU ETS. I abstract from any indirect effects, e.g. through equilibrium effects or the increase of electricity prices due to the EU ETS.

Future research could tackle some of these issues. For example, one could investigate the effects of the EU ETS on industry dynamics by employing structural approaches to ideally European firm-level data.

2.8 Acknowledgements

I thank Andreas Löschel, Ulrich Wagner, Kathrine von Graevenitz, Claudio Baccianti, Maximilian Auffhammer and Sofia Berto Villas-Boas for suggestions and insightful comments. I am indebted to the Research Data Centre (FDZ) of the Federal Statistical Office and the Statistical Offices of the German Länder for granting me access to the AFiD data and for the use of their research facilities, in particular Michael Rößner for his advice and technical support at different stages of this research. I gratefully acknowledge the financial support of the Helmholtz Association through the Helmholtz Alliance ENERGY-TRANS and the Helmholtz scholarship that enabled me to visit the Department of Agricultural and Resource Economics at the University of California, Berkeley, where this paper was written. Further, I would like to thank the Department of Agricultural and Resource Economics for its hospitality - especially my host during the stay, Maximilian Auffhammer. Thorsten Doherr and Tobias Krabel provided excellent research assistance. The views expressed in this paper are those of the author and do not necessarily represent those of the ZEW or the FDZ.

2.9 Appendices

Appendix 2.A. Additional information on data and descriptive statistics

2.A.1. EUA price development

Appendix A.1. provides additional information on the EUA price development. Figure 2.4 shows the price series of the ICE-ECX EUA front year futures with the closest maturity.

Figure 2.4: EUA prices 2005 - 2013 (in EUR)

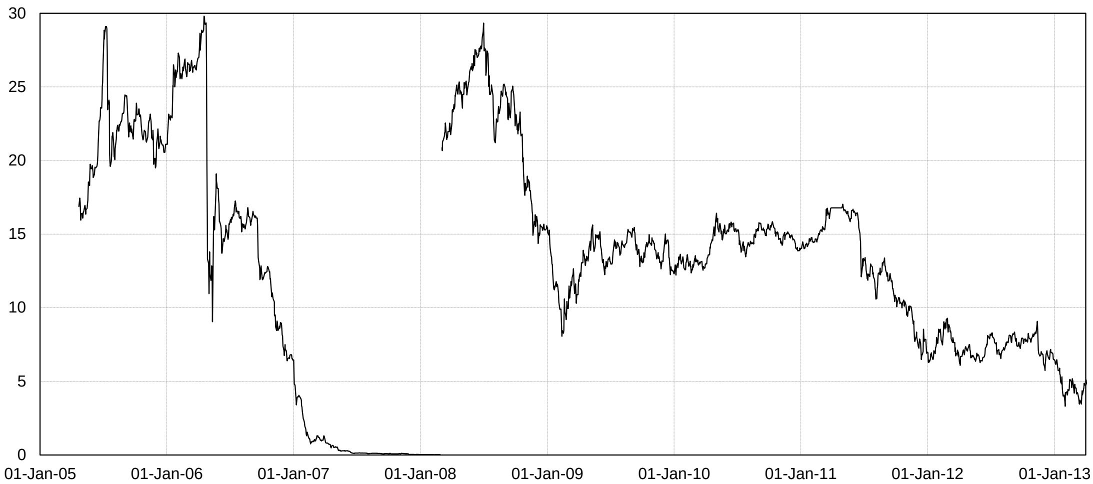

Notes: Source: ICE-ECX - access via Thomson Reuters Datastream.

2.A.2. Free allocation and verified emissions

Appendix A.2. provides information on the free allocation and the verified emissions of firms in the dataset. Table 2.8 shows the total amount of grandfathered EUAs and the total amount of verified emissions for two digit industries.

Table 2.8: Free allocation and verified emissions (in t CO_2)

		Phase I			Phase II		
NACE	Industry	Free allocation	Verified emissions	Balance	Free allocation	Verified emissions	Balance
10	Food products	10,730.32	9,695.56	1,034.76	21,471.65	18037.63	3434.02
11	Beverages	448.38	446.22	2.16	862.98	718.35	144.63
12	Tobacco products	.	.	.	.	.	.
13	Textiles	.	.	.	.	.	.
14	Wearing apparel	0	0	0	0	0	0
15	Leather and related products	0	0	0	0	0	0
16	Wood and products of wood and cork	1,805.56	575.27	1,230.29	5,320.81	1,029.93	4,290.88
17	Paper and paper products	21,398.24	17,597.27	3,800.97	36,300.12	26,661.00	9,638.12
18	Printing and reproduction of recorded media	.	.	.	.	.	.
19	Coke and refined petroleum products	77,947.19	77,925.31	21.88	134,922.3	124,959.69	9,962.61
20	Chemicals and chemical products	38,861.04	30,589.69	8,271.35	10,1483.09	87,533.43	13,949.66
21	Pharmaceutical products	2,276.38	1,390.66	885.73	3,108.41	2,040.37	1,068.03
22	Rubber and plastic products	1,130.28	1,000.44	129.85	1,775.60	1,463.27	312.34
23	Other non-metallic mineral products	90,831.29	79,230.95	11,600.34	156,617.63	137,016.66	19,600.97
24	Basic metals	112,048.83	101,291.07	10,757.76	264,550.73	164,814.81	99,735.92
25	Fabricated metal products	.	.	.	.	.	.
26	Computer, electronic and optical products	.	.	.	.	.	.
27	Electrical equipment	.	.	.	.	.	.
28	Machinery and equipment n.e.c.	.	.	.	.	.	.
29	Motor vehicles, trailers, and semi-trailers	8,853.423	8,154.973	698.45	12,476.513	14,623.63	-2,147.117
30	Other transport equipment	.	.	.	.	.	.
31	Furniture	0	0	0	0	0	0
32	Other manufacturing	.	.	.	.	.	.
33	Repair and installation of mach. and equip.	.	.	.	.	.	.

Notes: **.** marks statistics that have not been released by the statistical offices due to the low number of observations. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - Cost Structure Survey and AFiD-Panel Industrial Units, own calculations.

2.A.3. Matching AFiD, CSS, and EUTL

The different internal data sets of the Statistical Offices Germany, such as AFiD and CSS, can be easily merged via plant and firm-level identifiers. However, it requires some effort to match external data to AFiD and CSS, since the information on firm identifiers and names is not accessible for researchers. I match AFiD data on the firm-level with aggregated firm-level data from the EU Transaction Log for the years from 2005 to 2012 using the commercial register number and the value added tax number. I am able to match 77 percent (813 firms) of the firms in the EUTL with AFiD. The 238 firms that are not matched mainly belong to the energy, public, or service sector and thus are not contained in the production census for manufacturing.

2.A.4. Descriptive statistics - full sample

Table 2.9 reports additional descriptive statistics on the full sample.

Table 2.9: Additional descriptive statistics

	Mean	SD	Skewn	Kurtosis	p10	p50	p90	N
2000								
Gross value added	24,969.77	233,803.09	34.88	1,436.37	1,111.46	4,473.50	36,967.21	15,107
Output	58,790.54	548,088.98	48.64	3,017.09	2,077.68	10,260.49	96,753.55	14,998
Capital stock	21,722.82	176,689.16	33.88	1,380.38	455.93	3,545.52	34,412.70	15,039
Wages and salaries	12,995.41	116,283.62	47.35	2,665.08	797.66	2,868.15	22,520.11	15,106
Energy use	46,893.63	803,001.62	60.14	4,767.93	164.10	1,660.84	30,951.67	15,106
Number of employees	294.68	2,093.58	45.16	2,458.74	27.75	85.67	547.83	15,106
R&D expenditure	2,392.66	52,273.32	48.67	2,686.61	0	0	975.5043	15,107
Exports	30,226.70	459,974.70	49.82	2,919.78	0	1,111.121	35,350.72	15,106
2005								
Gross value added	27,345.33	266,307.21	39.93	2,011.60	1,173.94	4,987.07	39,760.41	13,858
Output	68,637.96	646,894.49	49.06	2,956.55	2,287.86	11,930.16	110,994.73	13,604
Capital stock	23,445.80	198,335.76	38.00	1,769.07	480.71	3,605.26	36,694.26	13,772
Wages and salaries	13,293.17	118,987.17	46.27	2,522.02	771.88	2,957.47	22,089.02	13,857
Energy use	55,565.39	793,966.15	42.11	2,313.17	275.40	2,135.74	36,675.84	13,669
Number of employees	296.61	2,106.77	44.87	2,419.63	28.25	88.92	531.25	13,857
R&D expenditure	3,082.25	62,091.20	44.90	2,293.96	0	0	1,541.87	13,858
Exports	39,977.06	560,038.10	46.71	2,517.40	0	2,106.06	48,757.09	13,857
2010								
Gross value added	25,483.97	269,640.70	39.65	1,836.76	1,149.20	4,582.03	34,591.33	15,399
Output	66,443.37	709,260.71	57.04	3,867.12	2,309.90	11,687.90	106,083.33	15,250
Capital stock	21,201.18	179,555.17	39.14	1,909.50	451.18	3,404.83	33,839.90	15,324
Wages and salaries	11,785.77	105,233.78	47.18	2,613.07	749.72	2,679.64	19,384.89	15,385
Energy use	62,728.99	1,144,134.97	51.02	3,124.93	294.94	2,136.16	35,485.39	15,258
Number of employees	269.29	1,887.62	46.79	2,628.07	30.00	85.08	479.58	15,384
R&D expenditure	2,878.184	60,147.87	46.24	2,480.41	0	0	1,389.11	15,399
Exports	40,749.75	639149.50	50.55	2918.14	0	2,050.15	48,962.84	15,385

Notes: Gross value added, output (production value), wages and salaries, R&D expenditure, exports and capital stock are denoted in EUR 1,000. Energy use is denoted in MWh. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - Cost Structure Survey, AFiD-Panel Industrial Units, and AFiD-Module Use of Energy, own calculations.

Appendix 2.B. Capital stocks for the German production census

Information on capital stocks is an important ingredient for several applications in empirical economic research - especially productivity analysis. The official production census of the German manufacturing sector (Amtliche Firmendaten für Deutschland; AFiD) comprises rich information on investments on the plant and the firm-level, but does not include information on capital stocks. In order to remedy this shortcoming, I compute capital stocks employing the perpetual inventory method (PIM).
The basic formula of the perpetual inventory method is

$$K_t = K_{t-1}(1-\delta) + I_t, \tag{2.13}$$

where K denotes capital stock, δ the geometric depreciation rate and I the investment. From the basic formula one can derive the initial capital stock K_1.

$$K_1 = I_0 + I_{-1}(1-\delta) + I_{-2}(1-\delta)^2 + ... \tag{2.14}$$

$$K_1 = \sum_{s=0}^{\infty} I_s(1-\delta)^s \tag{2.15}$$

I follow a notation that has been also used by earlier empirical studies as for instance Hall and Mairesse (1995). I assume the real investments to grow by the rate g.

$$K_1 = I_0 \sum_{s=0}^{\infty} \left[\frac{(1-\delta)}{(1+g)}\right]^s \tag{2.16}$$

$$K_1 = I_0 \frac{(1+g)}{(g+\delta)} \tag{2.17}$$

Hence, the capital at the beginning of the first period is defined by

$$K_1 = I_1 \frac{1}{(g+\delta)} \tag{2.18}$$

The PIM has been extensively used in studies that analyze sector and country level data. In principle, it also can be applied for the computation of capital stocks based on micro level data. However, turning from aggregate data toward micro data creates some issues that have to be considered: First, investments are lumpy, i.e. investments highly fluctuate over time. This property of investments at the plant and firm level creates difficulties to compute the initial capital stocks. Considering this, I compute the average of I_t over all time periods available in order to estimate I_1.

$$\hat{I_1} = \frac{\sum_{t=0}^{n} \frac{I_{t+1}}{(1+i)^t}}{n} \tag{2.19}$$

This leads us to the second issue that has to be considered when applying the PIM on micro data: It requires the observation of single agents (plants, firms, etc) over several time periods. For AFiD this issue is not a problem, since the investment data is census data. The resulting data set is a panel giving information for each year from 1995 to 2012. Otherwise it would be problematic to compute the initial capital stock. The growth rate of capital g and the depreciation rate δ can either be assumed to take a certain value or it can be estimated for each industry based on aggregate data. I follow the latter approach and use aggregated data at the two-digit industry level and compute industry specific average growth rates and depreciation rates.

In order to estimate the capital stock I use the firm-level investment data from the AFiD-Panel Industrial Units comprising investment in machinery and equipment, investment in buildings, and investment in property without buildings. I deflate the investments using two-digit industry-level deflators for machinery and equipment as well as general deflators for buildings and property without buildings. Starting from K_1 I plug the firm specific investments and the industry specific time-varying depreciation rates into Equation 2.13 in order to compute the entire time series of the firm's capital stock.

Apart from the AFiD-Panel Industrial Units, I exploit aggregate data at the two-digit industry level. These aggregate data can be retrieved via the Destatis portal GENESIS.[17] In particular I used the tables 81000-0107 National Accounts Depreciation, 81000-0115 Gross Investment, 81000-0116 Gross Capital Stock, 81000-0117 Net Capital Stock, and 61262-0001 Price Index Property in order to compute the growth rates, the depreciation rates and the deflators.

Since the focus of this study lies on the firm-level, I compute the capital stocks for firms. The method can be also employed to estimate capital stocks at the plant level. Table 2.10 reports descriptive statistics of the capital stock for firms in my sample.

[17] https://www-genesis.destatis.de/genesis/online

Table 2.10: Descriptive statistics capital stock

	Mean	SD	Skewn	Kurtosis	p10	p50	p90	N
1999	21,146.23	172,173.5	34.12	1,386.98	445.00	3,496.39	33,492.84	15,125
2000	21,722.82	176,689.2	33.88	1,380.38	455.93	3,545.52	34,412.70	15,039
2001	23,070.89	183,951.9	32.97	1,307.05	483.09	3,742.21	36,520.37	14,193
2002	23,910.89	190,909.2	32.59	1,269.70	498.21	3,867.37	37,574.70	13,603
2003	23,476.19	191,567.4	34.13	1,401.42	474.83	3,648.04	36,643.93	14,460
2004	23,071.02	191,259.8	37.10	1,684.88	472.18	3,574.86	36,549.22	14,270
2005	23,445.8	198,335.8	38.00	1,769.07	480.71	3,605.26	36,694.26	13,772
2006	23,777.36	197,148.1	37.34	1,719.16	484.33	3,666.73	37,776.86	13,405
2007	23,906.36	194,193.5	36.75	1,676.67	484.88	3,725.18	38,495.00	13,161
2008	20,665.96	174,540.6	40.26	2,021.59	450.68	3,340.48	33,742.33	16,088
2009	21,412.52	180,442.9	40.08	2,011.87	454.33	3,440.60	34,780.25	15,714
2010	21,201.18	179,555.2	39.14	1,909.50	451.18	3,404.83	33,839.90	15,324
2011	21,480.04	184,980.6	37.12	1,736.09	448.65	3,402.04	32,790.13	14,965
2012	21,393.36	189,500.1	37.90	1,795.70	433.42	3,146.24	33,135.09	14,926

Notes: Gross value added, output (production value), wages and salaries, R&D expenditure, exports and capital stock are denoted in EUR 1,000. Energy use is denoted in MWh. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - Cost Structure Survey, AFiD-Panel Industrial Units, and AFiD-Module Use of Energy, own calculations.

Chapter 3

The impact of the EU ETS on economic performance of German manufacturing firms

3.1 Introduction

The European Union (EU) aims at steering the European economy toward a competitive low-carbon pathway by 2050. Key to the EU's strategy is the EU Emissions Trading System (EU ETS) which was established in 2005 in order to cost-effectively curb greenhouse gas emissions from industrial installations. It is the world's largest international cap-and-trade system encompassing about 45 percent of the total European greenhouse gas emissions by regulating more than 11,000 installations in 31 countries. Despite the importance of the EU ETS, empirical evidence on its causal effects on the behavior of regulated firms is still scarce.

The EU ETS puts a price on greenhouse gas emissions from regulated installations and thus influences the production and investment decisions of regulated firms. There is concern that the EU ETS creates disadvantages for regulated firms exposed to competition from outside the EU ETS countries. In particular, firms from the manufacturing sector that sell their goods and services on global markets might be vulnerable due to additional cost imposed through the EU ETS. In this paper, we study the causal effect of the EU ETS on the economic performance of regulated firms from German manufacturing using official firm-level data.

For the evaluation of the EU ETS, the German manufacturing sector is a case of particular interest for two reasons. First, Germany is the largest economy and the largest emitter of greenhouse gas in the EU. In 2013, Germany emitted about 21 percent of the EU's total greenhouse gas emissions amounting to 976.3 million tonnes of CO_2 equivalent (Eurostat, 2016). Second, German manufacturing is export oriented and therefore may be especially affected by putting a price on carbon. In 2013, the German manufacturing sector exported almost 50 percent of the produced goods and services (Destatis, 2015).

A few very recent and mostly unpublished papers shed light on the causal effects of the EU ETS on regulated firms. Petrick and Wagner (2014) investigate the impact of the EU ETS on emissions, output, employment, and exports of manufacturing firms in Germany. They combine a difference-in-differences approach with semiparametric matching and weighted regressions in order to isolate the effect of the EU ETS. They show that the EU ETS reduced emissions of regulated firms by 20 percent during the years from 2007 to 2010. They do not find a significant negative effect of the EU ETS on employment, output, and exports. Following a similar approach, Wagner, Muûls, Martin, and Colmer (2014) show that the EU ETS reduced emissions of French manufacturing plants, by 15 to 20 percent on average between 2007 and 2010. They also find a significant decrease in employment in regulated plants of about 7 percent during the second compliance period of the EU ETS. Jaraitė and Di Maria (2016) investigate the impact of the EU ETS on Lithuanian firms employing nearest neighbor and kernel matching. They find that the EU ETS did not reduce CO_2 emissions, but improved CO_2 intensity. They do not find a significant effect on profits. However, regulated firms in Lithuania retired parts of their less efficient capital stock and made additional investments in the end of the second compliance period. Klemetsen, Rosendahl, and Jakobsen (2016) use a parametric difference-in-differences approach in order to isolate and quantify the effect of the EU ETS on emissions, emission intensity, value added, and labor productivity of Norwegian plants. They find that the EU ETS decreased emissions and at the same time increased value added and labor productivity during the second compliance period. Calel and Dechezleprêtre (2016) examine the effect of the EU ETS on on technological change, in particular patenting. They combine patent and commercial firm-level data for Europe with data from the EU ETS. Using a matching approach, they find that the EU ETS increased the number of low-carbon patents among regulated firms by 10 percent between 2005 and 2010 while not crowding out patenting for other technologies. Lutz (2016) estimates a structural production function that allows for endogenous productivity and employs a parametric difference-in-differences approach in order to quantify the effect of the EU ETS on firm-level productivity. He shows that the EU ETS had a significant positive impact on productivity during the first compliance period. [1]

We aim to contribute to the literature investigating the effect of the EU ETS on the economic performance of German manufacturing firms. So far, the literature examines the impact of emissions trading on output and the use of inputs separately or assesses firm performance relative to the mean production function of an industry. In contrast, we use a measure of economic performance that relates input use and produced output and assesses performance relative to the most efficient firms of the industry: We estimate firm-level technical efficiencies based on the stochastic production frontier approach by Aigner, Lovell, and Schmidt (1977).[2] The technical efficiency can be interpreted as distance to the frontier of the production set that is determined by the efficient firms of the corresponding industry. Subsequently, we employ different identification strategies in order to identify and estimate the effect of the EU ETS on the technical efficiency of regulated firms.

Following the studies depicted above, we exploit the installation-level inclusion criteria of the EU ETS that create variation in treatment. The EU ETS only covers emissions of installations with a capacity that exceeds thresholds determined by legislation. As a consequence, only firms operating large installations are covered by the EU ETS. The inclusion criteria open the avenue for identifying the effects of the EU ETS based on an array of suitable identification strategies. We use a difference-in-differences framework combined with an array of parametric conditioning strategies and nearest neighbor matching in order to identify and estimate the effect of the EU ETS on the economic performance of the regulated firms.

We use official firm-level data that is collected by the German statistical offices. It comprises general characteristics, such as revenues, value added, employment, and investment and is particularly detailed with regard to fuel and electricity use. The data serves as a basis for many official German governmental statistics and includes all manufacturing firms with more than 20 employees. Our panel covers two pretreatment years (2003-2004)

[1] Our review of the recent literature focuses on studies that aim to investigate the causal effect of the EU ETS on regulated firms. For a comprehensive overview of the literature on the EU ETS, we refer to Martin, Muûls, and Wagner (2016)

[2] This approach has been used in several studies that evaluate regulatory intervention, such as Knittel (2002).

as well as the first (2005-2007) and the second compliance period (2008-2012) of the EU ETS.

Applying a difference-in-differences approach combined with parametric conditioning strategies to the full census, we do not find a significant effect of the EU ETS on the performance of regulated firms. In order to investigate potential heterogeneous treatment effects across industries, we conduct a subsample analysis following the same design. We estimate the treatment effect on the two-digit level for the industries manufacture of food products (10), manufacture of paper and paper products (17), manufacture of chemicals and chemical products (20), and manufacture non-metallic mineral products (23). We find that some industries remain unaffected, while others display economically and statistically significant impact of the EU ETS on technical efficiency. The EU ETS had a significant negative impact on the technical efficiency of the regulated firms. In other words, the EU ETS decreased the average distance to the frontier or increased the economic performance of regulated firms.

In addition, we employ nearest neighbor matching to account for observable differences between treated and untreated firms. This alternative identification strategies allows us to relax the parametric assumptions of the standard difference-and-differences approach that are applied to the treatment and outcome model. The results of nearest neighbor matching indicate a statistically and economically significant negative effect of the EU ETS on the technical efficiency of the regulated firms during the first compliance period.

The remainder of our paper is organized as follows. In Section 3.2, we describe the regulatory framework of the EU ETS. Section 3.3 outlines the identification strategy employed to isolate the effect of the EU ETS on technical efficiency. Section 3.4 describes the German production census and additional data sources. Section 3.5 reports the results and Section 3.6 concludes.

3.2 The EU ETS

The EU ETS is the largest multinational cap-and-trade system covering around 45 percent of the EU's greenhouse gas emissions. As core instrument of EU climate policy, it was enacted by Directive 2003/87/EC in October 2003 and finally implemented in January 2005 (European Parliament and Council, 2003). The EU ETS regulates the emissions of more than 11,000 energy-intensive industrial installations across the 31 countries of the European Economic Area (EEA)[3].

The EU ETS is organized in temporally separated compliance periods. Phase I (2005 - 2007) is marked as pilot or introductory phase. Since only few member states had experiences with emissions trading, the European Commission accorded regulators and firms time to adapt to this new instrument.[4] Phase II (2008 - 2012) of the EU ETS corresponds to the commitment period of the Kyoto Protocol and was linked to the flexible mechanisms that certified project-based emission reductions, the Clean Development Mechanism (CDM) and Joint Implementation (JI) (European Parliament and Council, 2004). The following Phases III (2013 - 2020) and IV (2021 - 2030) are linked to the emission targets in the 2020 Climate and Energy Package and the 2030 Climate and Energy Framework, respectively (European Parliament and Council, 2009).

According to the milestones of the EU climate policy, the cap of the EU ETS is annually lowered by 1.74 percent during Phase III. This corresponds a reduction of emissions by 21 percent relative to 2005 in 2020. From 2021 onwards, the cap will be anually decreased by 2.2 percent (European Council, 2014). The emission rights that are traded in the framework of the EU ETS are referred to as European Union Allowances (EUAs). One EUA corresponds to one metric tonne of CO_2 equivalent. Each year, firms that are regulated by the EU ETS must surrender EUAs according to their verified emissions. Since January 2012, the EU organizes an allowance market specifically for aviation operators active in the EEA.

During the first two compliance periods, the main mode of allocation was grandfathering. The allocation of allowances was governed decentralized at the member state level by the National Allocation Plans. Furthermore, member states were responsible for setting

[3]The EEA includes the 28 EU member states as well as in Iceland, Liechtenstein, and Norway.

[4]Only UK and Denmark had experiences with national greenhouse gas emissions trading systems when the EU ETS was established in 2005.

up national registries to record the issuance, transfer, and surrender of EUAs. The EC supervised the national emission registries by maintaining the Community Independent Transaction Log (CITL). Emissions of regulated installations are monitored and reported annually by the firm and verified by independent referees. The penalty for non-compliance with the EU ETS was 40 euros per EUA in Phase I and 100 euros in Phase II. From the beginning of Phase III, the allowance allocation was centralized and the main mode of allocation started to gradually shift from grandfathering to auctioning.

Our analysis focuses on the first two compliance periods of the EU ETS. Phase I was completely decoupled from Phase II. Banking and borrowing was allowed across years within each compliance period, but not between Phase I and II. As a consequence, a tremendous over-allocation of free EUAs during Phase I led to a decline in EUA prices from above 25 euros to zero in 2007. In Phase II, the EU ETS also suffered from massive over-allocation. Due to the decline in economic activity and thus CO_2 emissions in the wake of the economic crisis, the unadjusted supply of free allowances led to an oversupply of allowances. This development was enhanced by the heavy use of certificates issued by CDM and JI projects. In contrast to Phase I, however, it was possible to bank EUAs for future use in the following compliance periods. As a result of these developments, the EUA price decreased from more than 25 euros at the beginning of Phase II to less than 10 euros in the second half of Phase II.[5]

In the manufacturing sector, combustion installations for the generation of electric power and heat with a rated thermal input in excess of 20 megawatts as well as energy intensive production processes are regulated. These processes include oil refining, the production and processing of ferrous metals, the manufacture of cement, the manufacture of lime, ceramics including bricks, glass, and the production and processing of pulp and paper are regulated. The EU ETS only regulates large installations with capacities in excess of process-specific thresholds determined by regulation.[6] Table 3.1 shows the total number of firms and the number of regulated firms in our data set of the German manufacturing sector across two-digit industries classified by the NACE code. The regulated processes are concentrated in a few energy intensive industries.

[5] More details on the EUA price development can be found in Appendix 3.B.

[6] More details on the inclusion criteria of the EU ETS can be found in European Parliament and Council, 2003).

There exist firms both regulated and unregulated in the same industries. The inclusion criteria therefore create variation in the treatment status and enable us to identify the causal effects of the EU ETS. We will take into account the structural differences across regulated and unregulated firms by using different parametric and nonparametric strategies explained in the following section.

Table 3.1: Number of observations by industry - German production census

		2005		2008		2012	
NACE	Industry	Total	Regulated	Total	Regulated	Total	Regulated
10	Food products	4,653	50	4,680	53	4,831	54
11	Beverages	601	11	534	13	483	15
12	Tobacco products	23	1	22	2	21	2
13	Textiles	809	7	734	7	654	7
14	Wearing apparel	470	-	383	-	277	0
15	Leather and related products	180	-	160	-	123	0
16	Wood and products of wood and cork	1,317	14	1,195	21	1,124	19
17	Paper and paper products	829	89	809	97	789	100
18	Printing and reproduction of recorded media	1,608	2	1,543	2	1,335	3
19	Coke and refined petroleum products	48	16	48	17	47	16
20	Chemicals and chemical products	1,140	56	1,166	58	1,194	55
21	Pharmaceutical products	273	8	261	8	255	7
22	Rubber and plastic products	2,698	12	2,730	12	2,765	14
23	Other non-metallic mineral products	1,789	162	1,635	159	1,570	155
24	Basic metals	903	33	923	34	915	35
25	Fabricated metal products	6,111	3	6,410	5	6,820	4
26	Computer, electronic and optical products	1,677	5	1,687	4	1,637	4
27	Electrical equipment	1,975	5	2,015	5	1,914	5
28	Machinery and equipment n.e.c.	5,919	6	6,134	8	5,296	8
29	Motor vehicles, trailers, and semi-trailers	1,127	10	1,130	9	1,015	9
30	Other transport equipment	319	5	329	5	251	5
31	Furniture	1,041	-	1,005	-	971	-
32	Other manufacturing	1,560	3	1,472	3	1,432	2
33	Repair and installation of machinery and equipment	295	-	289	-	1,482	1
–	Total	37,365	498	37,294	522	37,201	520

Notes: Number of firms for the first year of Phase I of the EU ETS (2005), the first year of Phase II (2008) and the last year of Phase II (2012) that is also the last year we observe. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

3.3 Empirical strategy

The measure of economic performance analyzed in this study is based on the technical efficiency of the firm. We rely on the stochastic production frontier concept introduced by Aigner, Lovell, and Schmidt (1977) and Meeusen and van den Broeck (1977) in order to quantify the technical efficiency. In contrast to the estimation of average production functions, the stochastic frontier analysis enables the estimation of a production frontier. This function expresses the maximum amount of output that can be produced from a given set of inputs with a fixed technology. We use the firm specific distance to the frontier - the technical efficiency - as measure of economic performance.[7]

In order to account for industry specific technologies, we estimate the stochastic frontier model for each two-digit industry within the manufacturing sector. The implied technical efficiency refers to a frontier that is estimated using data from 2003 to 2012. In this way, the estimated technical efficiency also captures dynamic factors that might drive firm's efficiency, such as technological change. Our identification strategy will take these characteristics of the technical efficiency into account.

3.3.1 Identifying the effect of the EU ETS

The EU ETS only covers CO_2 emissions of installations with a capacity that exceeds thresholds determined by the regulatory authorities.[8] We exploit this variation created by the inclusion criteria of the EU ETS in order to isolate the effect of the EU ETS on the technical efficiency of regulated firms. We follow the literature on program evaluation and employ the potential outcome framework introduced by Rubin (1974, 1977).[9] We differentiate between treatment and control group depending on whether a firm has to comply with the regulation by the EU ETS or not. Let the binary variable $ETS_i \in \{0, 1\}$ be an indicator that describes the treatment status of firm i. Let ETS_i be equal 1 if the firm operates installations that are regulated by the EU ETS and 0 if the firm is not

[7] In Appendix 3.C, we describe the estimation of the production frontier and the recovery of the technical efficiencies in more detail.

[8] See Section 3.2 for details.

[9] The potential outcome framework has become a common way to describe an identification strategy in policy evaluation literature. Also, studies investigating the effects of emission trading schemes frequently rely on the potential outcome framework, see for instance Fowlie, Holland, and Mansur (2012).

required to participate in the EU ETS. Accordingly, we describe the potential outcomes by $Y_i(1)$ and $Y_i(0)$ for treatment and control group, respectively. Our aim is to estimate the sample average treatment effect on the treated (SATT):

$$\tau = E[Y_{it}(1) - Y_{it}(0)|ETS_i = 1], \tag{3.1}$$

where τ is the average effect of the EU ETS on the technical efficiency of regulated firms after the implementation of the EU ETS. While we are able to observe $Y_{it}(1)$ for regulated firms, the outcome $Y_{it}(0)$ is not realized in the case of regulated firms. Therefore, we will use information on the outcome $Y_{it}(0)$ collected from the firms that belong to the control group in order to form an adequate counterfactual. The comparison of the two groups will only lead to robust results, if factors that are correlated with technical efficiency dynamics do not differ across treatment and control group. In the following sections, we will present strategies that take this potential source of bias into account.

3.3.2 Difference-in-differences

We start from a baseline difference-in-differences specification. In order to control for observed and constant unobserved confounding factors, we gradually enhance the model by including explanatory variables and firm-level fixed effects into the estimation equation.[10]

The key identifying assumption of our baseline difference-in-differences specification is, that the technical efficiency trends would be the same in the treatment and control group in the absence of the EU ETS. We will investigate the validity of the common trend assumption by analyzing pretreatment developments of technical efficiency across treatment and control group in Section 3.5.2. In addition, we assume that the EU ETS only has an effect on regulated entities. This assumption is often referred to as stable unit treatment value assumption (SUTVA) and basically excludes general equilibrium effects and spillover effects.

Our baseline specification of the difference-in-differences model takes the form

$$\ln Y_{it} = \beta_0 + \beta_1 ETS_i + \tau\ ETS_i \times \mathrm{I}(t \geq 2005) + \varphi_t + \gamma_s + \eta_{st} + \varepsilon_{it}, \tag{3.2}$$

[10] The procedure to start from a baseline difference-in-differences approach and then to enhance it gradually is quite common in the program evaluation literature. Gray, Shadbegian, Wang, and Meral (2014) employ a similar approach to investigate the effects of environmental regulation on employment of the U.S. paper industry. Lutz (2016) and Klemetsen, Rosendahl, and Jakobsen (2016) choose this strategy in order to identify the effect of the EU ETS on German firms and Norwegian plants, respectively.

where Y_{it} denotes the outcome variable technical efficiency of firm i in year t, as described above. ETS_i is a dummy that indicates if the firm must comply with the EU ETS, $I(t \geq 2005)$ is a dummy that indicates if the year t lies in the treatment period, φ_t is a year fixed effect, γ_s is an industry fixed effect, η_{st} is the interaction term of year and industry fixed effect, and ε_{it} is a zero mean error term. Our interest lies in the coefficient τ that measures the average treatment effect of the EU ETS on the technical efficiency of the regulated firms.

For our baseline specification, we assume that the counterfactual technical efficiency is equally distributed across treatment and control group conditional on group, two-digit industry, and year fixed effects and a full set of interaction terms. We relax this conditional unconfoundedness assumption by controlling for additional confounding factors that might be correlated with both the treatment and the technical efficiency. Since the compliance with the EU ETS depends on the capacity of the installation, especially factors related to the scale of the production and the size of the firm might impede the estimation of the average treatment effect. Regrettably, we do not observe the capacity, but we include among other controls the value of the physical capital stock in order to take scale effects into account. We consider the following specification of the difference-in-differences model that includes additional explanatory variables:

$$\ln Y_{it} = \beta_0 + \beta_1 ETS_i + \tau\ ETS_i \times I(t \geq 2005) + \mathbf{z}_{it}\Psi + \varphi_t + \gamma_s + \eta_{st} + \varepsilon_{it}, \qquad (3.3)$$

where $\mathbf{z}_{it}$ denotes a vector of firm characteristics and Ψ is the vector with the corresponding coefficients. We now further relax the assumption of conditional unconfoundedness by allowing for unobserved constant factors. In particular, we estimate a specification of the difference-in-differences model that includes a firm-level fixed effect:

$$\ln Y_{it} = \beta_1 ETS_i + \tau\ ETS_i \times I(t \geq 2005) + \mathbf{z}_{it}\Psi + \alpha_i + \varphi_t + \eta_{st} + \varepsilon_{it}, \qquad (3.4)$$

where α_i denotes the firm-level fixed effect that captures constant characteristics of the firm, such as average capacity and location. [11]

[11] Industry fixed effects drop out, since these are constant over time.

3.3.3 Nearest neighbor matching

In addition to the parametric difference-in-differences model, we estimate a model based on nearest neighbor matching in order to relax the assumptions on the functional form of the treatment and outcome model. In the program evaluation literature on emission trading schemes, matching has become quite popular in recent years. Fowlie, Holland, and Mansur (2012) employ a nonparametric matching strategy in order investigate the effectivity of the Californian RECLAIM program. Petrick and Wagner (2014), Wagner, Muuls, Martin, and Colmer(2014), Jaraitė and Di Maria (2016), and Calel and Dechezlepretre (2016) implement different matching approaches in order to investigate the impact of the EU ETS on emissions, competitiveness, and R&D activities of regulated firms. Our matching approach is closely related to the one employed by Fowlie, Holland, and Mansur (2012), since we use nonparametric nearest neighbor matching in order to form an adequate control group. The matching approach enables us to relax some of the assumptions we have to make in the framework of the difference-in-differences approach described above. We do not pose any parametric assumptions on the relationship between technical efficiency and the explanatory variables $\mathbf{z}_{it}$. However, we still rely on the conditional unconfoundedness and SUTVA. For the matching approach, the common support assumption is of particular importance, i.e. we assume the conditional probability to be treated is larger than 0 and smaller than 1: $0 < P[ETS_i = 1|X] < 1$. In accordance with Heckman, Ichimura, and Todd (1997) and Heckman, Ichimura, Smith, and Todd (1998), we estimate the average treatment effect using the following difference-in-differences matching estimator

$$\hat{\tau} = \frac{1}{N} \sum_{j \in I_1} \Big\{ (Y_{jt'}(1) - Y_{jt^0}(0)) - \sum_{k \in I_0} w_{jk} (Y_{kt'}(0) - Y_{kt^0}(0)) \Big\}, \tag{3.5}$$

where I_1 denotes the set of regulated firms, I_0 denotes the set of the unregulated firms, N is the number of firms in the treatment group. The regulated firms are indexed by j, whereas the unregulated firms are indexed by k. Let w_{jk} denote the weight placed on firm k when constructing the counterfactual estimated for the treated firms. We employ matching on firm characteristics within two-digit industries in order to form an adequate control group.

3.4 Data

We employ official firm-level data collected by the German Federal Statistical Office and the Statistical Offices of the German Federal States. The Official Firm Data for Germany (Amtliche Firmendaten für Deutschland - AFiD) is a highly reliable data source that forms the basis of many official German governmental statistics. The participation in the underlying production census is mandatory by law and the results of the conducted surveys are validated by the statistical offices.

We have remote access to annual data from 2003 to 2012.[12] AFiD is of modular nature, i.e. the statistical offices conduct annual surveys on different topics and combine the collected data to thematic modules. We use the longitudinal census database AFiD-Panel Industrial Units that contains annual data from the Monthly Report on Plant Operation, the Census on Production, and the Census on Investment. This module contains detailed information on inputs and outputs that describe the production process. In addition, we use the AFiD-Module Use of Energy. It is a longitudinal census that combines results from the Census on Energy Use and the Monthly Report on Plant Operation. It includes comprehensive data on electricity and fuel purchase, sale, and use. The AFiD-Panel Industrial Units and the AFiD-Module Use of Energy have the same group of respondents: All German plants that are active in manufacturing and belong to firms that employ more than 20 persons must participate in the underlying surveys. We aggregate plant-level data to the firm level using the firm affiliation provided by the AFiD-Panel Industrial Units. The firms are classified according to ISIC rev. 4.[13]

As output variable for our stochastic production frontier model, we employ the value of production in the corresponding year denoted in EUR. The output variable has been deflated using two-digit industry specific price indices. The capital stock is computed by applying the perpetual inventory method to the investment data contained in the AFiD-Panel Industrial Units and is denoted in EUR. A detailed description of the methodology and its application to AFiD data can be found in Lutz (2016). The number of employees in the firm indicates the use of labor. The aggregated energy use is computed based on the

[12] We also have access to data for the years from 1995 to 2002. However, the statistical offices changed the survey gathering the information on energy use in 2003 making it difficult to include the data before 2003 into our investigation.

[13] In Appendix 3.A, we present more information on the industry classification.

electricity and fuel use information contained in the AFiD-Module Use of Energy Use and is measured in MWh. We compute the CO_2 emissions from the fuel use and the net use of electricity contained in the AFiD-Module Use of energy using data on CO_2 content in fuels and electricity from the German statistical offices and the Federal Environmental Agency.[14] The computation of the emissions as well as the emission coefficients are described in Appendix 3.A. The CO_2 emissions are measured in t CO_2 equivalent.

In order to identify firms that are regulated by the EU ETS, we match the production census with data of the European Union Transaction Log (EUTL) from the years 2005 to 2012 using the commercial register number and the VAT number. During the period from 2005 to 2012, a total of 1051 German firms was regulated by the EU ETS. We are able to match 77 percent (813 firms) of the firms in the EUTL with AFiD. The remaining 238 firms mainly belong to the energy sector, the public sector (hospitals and universities), or the service sector (e.g. airports and exhibition centers) and thus could not be matched with a production census that only contains information of manufacturing firms.

Table 3.2 shows descriptive statistics of the variables used in the study for the entire manufacturing sector. The output as well as the use of inputs increase over time. However, the economic crisis is reflected in the descriptive statistics for the year 2009. In particular, output, emissions, and energy use declined. The number of employees remained quite stable due to the support programs and the strict labor market regulation. The capital stock also remained quite stable, however it has slightly decreased in the aftermath of the crisis due to low investments during the crisis. The number of observations vary across variables within years, since the information is collected through different surveys as explained above.[15]

Figure 3.1 sheds some light onto the development of the employed variables over time across two-digit industries within manufacturing. We plot the development of the indexed median (base year 2003) of each variable for the food industry (10), the paper industry (17), the chemical industry (20), and manufacture of non-metallic mineral products (23, comprises for example the glass and cement industry).[16] These four sectors cover more than half of the German manufacturing firms regulated by the EU ETS. While the de-

[14]We describe the computation of the CO_2 emissions in Appendix 3.A.

[15]The surveys are not conducted at the exact same date and thus the number of firms might vary to a minor degree.

[16]Appendix 3.B contains the graphs for each two-digit industry in manufacturing.

velopment of output as well as input use in the food industry was barely affected by the economic crisis, the other graphs for these industries show a strong impact on output, emissions, and energy use in 2009. As we will learn in Section 3.5.1, this will be also reflected in the technical efficiencies, since firms produced less in the crisis year while they were not able to adjust their capital stock and their use of labor in the short term. The former has a strong dynamic character and can only be adjusted through investment or the disposal of physical capital. The latter also has a dynamic character in Germany. Due to strong labor market regulation and collective labor agreements, long periods of notice prevent short-term adaption of the labor force. Figure 3.1 also suggest a strong relation between output and energy use.

Table 3.2: Descriptive statistics German production census

	Mean	SD	Skewness	Kurtosis	P10	P50	P90	N
2003								
Output (in EUR 1,000)	28,699.22	360,632.64	81.42	8,262.81	1,324.71	5,276.71	42,903.50	37,888
Emissions (in t CO_2)	7,343.68	119,327.97	49.20	2,996.19	64.39	374.07	5,236.04	36,985
Capitalstock (in EUR 1,000)	11,322.21	120,382.53	53.92	3,524.23	256.28	1,838.52	16,220.64	37,099
Number of employees	153.06	1,292.85	71.69	6,278.85	22.58	50.00	254.00	38,319
Energy use (in MWh)	21,418.67	376,270.67	49.81	2,991.01	161.79	911.97	12,979.47	36,949
2006								
Output (in EUR 1,000)	33,903.31	417,111.75	77.80	7,509.40	1,483.30	6,201.49	49,867.54	36,162
Emissions (in t CO_2)	8,978.83	193,026.98	69.25	6,186.59	72.56	412.27	5,742.81	35,654
Capitalstock (in EUR 1,000)	11,097.78	121,072.64	60.25	4,516.43	248.68	1,774.37	15,880.93	36,073
Number of employees	154.01	1,313.92	73.58	6,446.77	23.90	52.92	254.75	36,632
Energy use (in MWh)	27,200.64	618,927.15	64.61	5,055.30	186.25	994.59	14,296.39	35,631
2009								
Output (in EUR 1,000)	29,257.35	345,810.11	77.27	7,309.36	1,295.66	5,346.61	44,534.20	36,703
Emissions (in t CO_2)	7,989.21	179,143.00	69.44	5,857.38	66.90	362.56	5,017.29	36,100
Capitalstock (in EUR 1,000)	11,148.64	119,355.75	60.18	4,565.63	234.60	1,785.08	16,464.80	36,335
Number of employees	152.47	1,219.86	70.98	6,127.83	24.00	53.00	254.50	36,982
Energy use (in MWh)	26,043.38	627,994.42	70.33	5,934.58	179.39	920.20	13,337.29	36,074
2012								
Output (in EUR 1,000)	35,194.48	514,872.08	89.04	9,350.33	1,431.54	6,184.66	51,415.54	36,882
Emissions (in t CO_2)	9,012.12	211,455.53	71.14	6,276.90	68.58	385.09	5,489.93	36,435
Capitalstock (in EUR 1,000)	10,641.92	122,387.63	58.11	4,256.77	240.51	1,611.20	14,924.90	36,380
Number of employees	157.09	1,300.76	69.85	5,901.93	25.00	54.83	260.33	37,130
Energy use (in MWh)	29,383.97	745,139.74	73.12	6,467.69	183.05	951.67	14,134.01	36,421

Notes: Output (production value), and capital stock are denoted in EUR 1,000. Energy use is denoted in MWh and CO_2 emissions in t CO_2. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

Figure 3.1: Descriptive statistics: development across industries

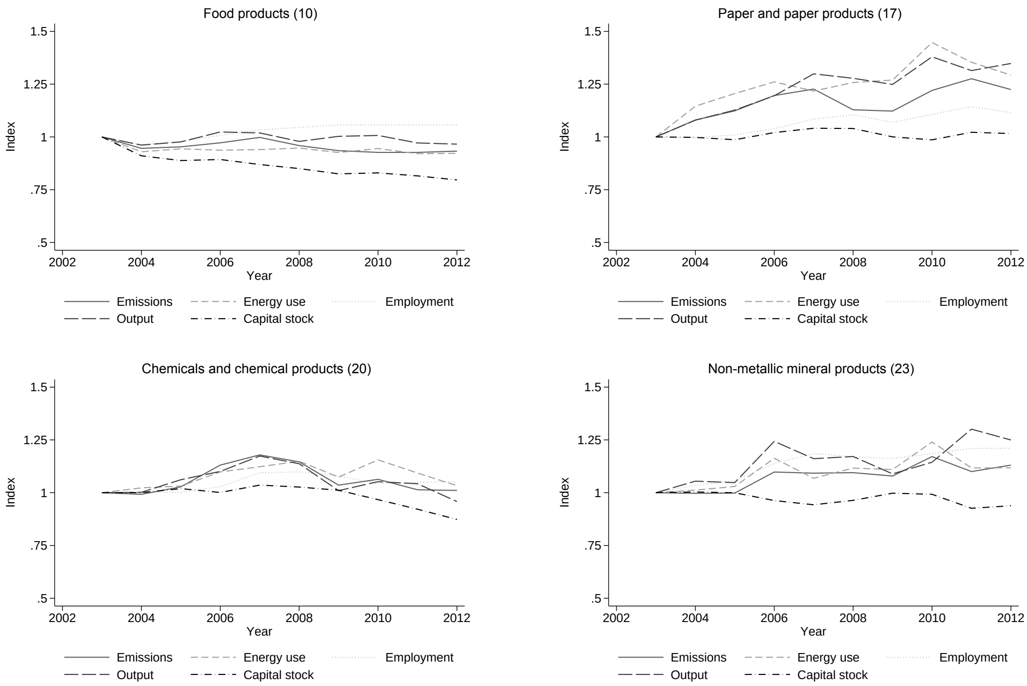

Notes: Indexed medians (base year 2003) for emissions, energy use, employment, output, and capital stock. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

3.5 Results

In this section, we present the parameter estimates of the stochastic production frontier model and shed light onto the development of technical efficiency across treatment and control group. We then empirically examine the core assumptions of our identification strategies and finally show the estimated treatment effects based on the different approaches.

3.5.1 Stochastic production frontiers and technical efficiencies

The industries within the manufacturing sector differ considerably with respect to produced goods, production processes, and market structures. We take this heterogeneity into account and estimate separate Cobb-Douglas production frontiers for each two-digit industry providing a common point of reference for the entire time period from 2003 to 2012. Table 3.3 shows the parameter estimates of the stochastic production frontier model.

The estimated parameters of the stochastic production frontier vary across industries reflecting the strong heterogeneity within the manufacturing sector. The economies of scale also vary across industries and range between 0.93 (manufacture of other non-metallic mineral products; 23) and 1.20 (manufacture of beverages; 11). For the majority of industries, we observe statistically significant increasing economies of scale. Table 3.3 also shows the parameter estimates that characterize the distributions of the composite error term. The parameter $\hat{\sigma}_u$ denotes the estimated standard deviation of the mean zero normal distribution of the noise component u_{it}. The parameters $\hat{\mu}_\nu$ and $\hat{\sigma}_\nu$ denote the estimated mean and standard deviation of the truncated normal distribution of the technical efficiency component. The estimates for $\hat{\mu}_\nu$ are comparatively large, since we estimate a joint frontier for the entire time span from 2003 to 2012. This is necessary in order to obtain a single point of reference that allows for comparisons across years. As a robustness check, we also estimate the stochastic production frontiers using a value added representation. The results are similar to the results of the gross output representation and are reported in Appendix 3.D.

Now we turn toward the development of technical efficiency over time and examine differences across treatment and control group. We focus on the four industries manufacture of food products (10), manufacture of paper and paper products (17), manufacture

Table 3.3: Parameter estimates: production frontier

Industry (NACE)	# Firms	Capital	Labor	Energy	Constant	$\hat{\sigma}_u$	$\hat{\mu}_\nu$	$\hat{\sigma}_\nu$
Food products (10)	6935	0.265 (0.010)	0.323 (0.016)	0.481 (0.014)	2.047 (0.042)	0.609 (0.010)	-443.728 (6.786)	11.969 (0.384)
Beverages (11)	703	0.223 (0.032)	0.725 (0.050)	0.257 (0.036)	2.252 (0.188)	0.549 (0.023)	-365.726 (119.853)	9.839 (2.725)
Textiles (13)	1103	0.199 (0.020)	0.738 (0.037)	0.117 (0.016)	3.652 (0.104)	0.507 (0.019)	-512.914 (19.645)	13.732 (0.832)
Leather and related products (15)	231	0.203 (0.050)	0.742 (0.081)	0.177 (0.045)	3.308 (0.251)	0.514 (0.041)	-908.894 (27.054)	24.299 (1.631)
Wood and products of wood and cork (16)	1587	0.186 (0.017)	0.794 (0.029)	0.146 (0.012)	3.507 (0.079)	0.498 (0.016)	-513.569 (16.601)	13.775 (0.552)
Paper and paper products (17)	1104	0.178 (0.021)	0.677 (0.031)	0.183 (0.012)	3.720 (0.089)	0.389 (0.017)	-360.647 (15.225)	9.668 (0.581)
Printing and reproduction of recorded media (18)	2255	0.115 (0.013)	0.689 (0.026)	0.250 (0.014)	3.580 (0.061)	0.367 (0.011)	-363.949 (73.232)	9.821 (1.284)
Chemicals and chemical products (20)	1722	0.205 (0.024)	0.596 (0.029)	0.173 (0.014)	4.372 (0.092)	0.522 (0.016)	-607.081 (40.582)	16.373 (0.883)
Rubber and plastic products (22)	3935	0.155 (0.011)	0.726 (0.017)	0.178 (0.010)	3.645 (0.047)	0.416 (0.008)	-385.313 (44.152)	10.408 (0.750)
Other non-metallic mineral products (23)	2446	0.206 (0.014)	0.612 (0.020)	0.111 (0.009)	4.229 (0.070)	0.501 (0.013)	-471.085 (5.073)	12.644 (0.427)
Basic metals (24)	1274	0.241 (0.024)	0.637 (0.040)	0.163 (0.019)	3.617 (0.096)	0.610 (0.019)	-300.333 (11.398)	8.107 (0.761)
Fabricated metal products (25)	9676	0.103 (0.006)	0.896 (0.011)	0.112 (0.006)	3.791 (0.030)	0.458 (0.006)	-372.983 (1.690)	10.107 (0.190)
Electrical equipment (27)	3077	0.170 (0.011)	0.834 (0.021)	0.071 (0.011)	4.088 (0.049)	0.449 (0.010)	-501.796 (6.310)	13.482 (0.360)
Machinery and equipment n.e.c. (28)	8620	0.071 (0.006)	1.066 (0.011)	0.027 (0.007)	4.092 (0.032)	0.453 (0.006)	-404.643 (2.223)	10.965 (0.204)
Motor vehicles, trailers, and semi-trailers (29)	1681	0.167 (0.017)	0.893 (0.029)	0.067 (0.019)	3.840 (0.072)	0.589 (0.020)	-405.271 (53.125)	10.924 (1.072)
Furniture (31)	1532	0.133 (0.013)	1.034 (0.026)	0.036 (0.016)	3.599 (0.071)	0.433 (0.014)	-409.503 (7.892)	11.030 (0.481)

Notes: The number of observations includes all firms that were active during the period from 2003 to 2012. We do not consider the industries manufacture of tobacco products (12), manufacture of wearing apparel (14), manufacture of pharmaceutical products (21), manufacture of computer, electronic and optical products (26), manufacture of other transport equipment (30), other manufacturing (32), and repair and installation of machinery and equipment (33). For more information see Appendix 3.C. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

of chemicals and chemical products (20), and manufacture non-metallic mineral products (23). These industries contain a sufficiently high number of regulated firms and thus enable adequate statistical inference.[17] Figure 3.2 consists of four graphs showing the development of the median technical efficiency over time within the four industries.

Since we estimate one stochastic frontier per industry that serves as reference point for the entire time period from 2003 to 2012, the dynamics of the technical efficiency reflect two developments. First, we observe that in all four industries, the median technical efficiency decreases during the early 2000s, i.e. the median firm becomes more efficient relative to the firms operating on the frontier. This negative trend in technical efficiency is driven by technical progress. We observe that the estimated stochastic frontier is determined by observations during the more recent years. Over time, the technical efficiency decreases, since technical progress gradually pushes the firms toward the frontier. Secondly, we observe increases in technical efficiency from 2006 onwards coinciding with the economic crisis. The technical efficiency peaks in 2009, the year when the crisis hit German manufacturing hardest. While demand and thus the production of goods rapidly decrease, firms do not adjust their capacity at the same speed. Therefore, low utilization rates increase the distance to the frontier during the economic crisis (see Section 3.4 for details on input use). Our empirical strategy is not impaired by these developments as long as treatment and control group are equally affected conditional on observable firm characteristics and an array of fixed effects that depend on the estimated specification.

The median technical efficiency of the treated firms is portrayed by the dotted line. It is higher than the median technical efficiency of the control group indicating that the treated median firm operates less efficiently in these industries. The dashed line displays the development of the median technical efficiency in the control group. It is close to the line of the overall median technical efficiency reflecting that the share of control firms is high. For the food industry (10), the paper industry (17), and the chemical industry (20), the distance between the median technical efficiency of treatment and control group decreases over time. In particular during the years from 2005 to 2007, the median technical efficiency of the treated firms declines and thus converges toward the median technical efficiency of the control group. The non-metallic mineral industry (23) does not show such a development.

[17] For the subsample analysis, we only consider two-digit industries with at least 50 regulated firms.

Figure 3.2: Comparison of treatment and control groups: technical efficiencies

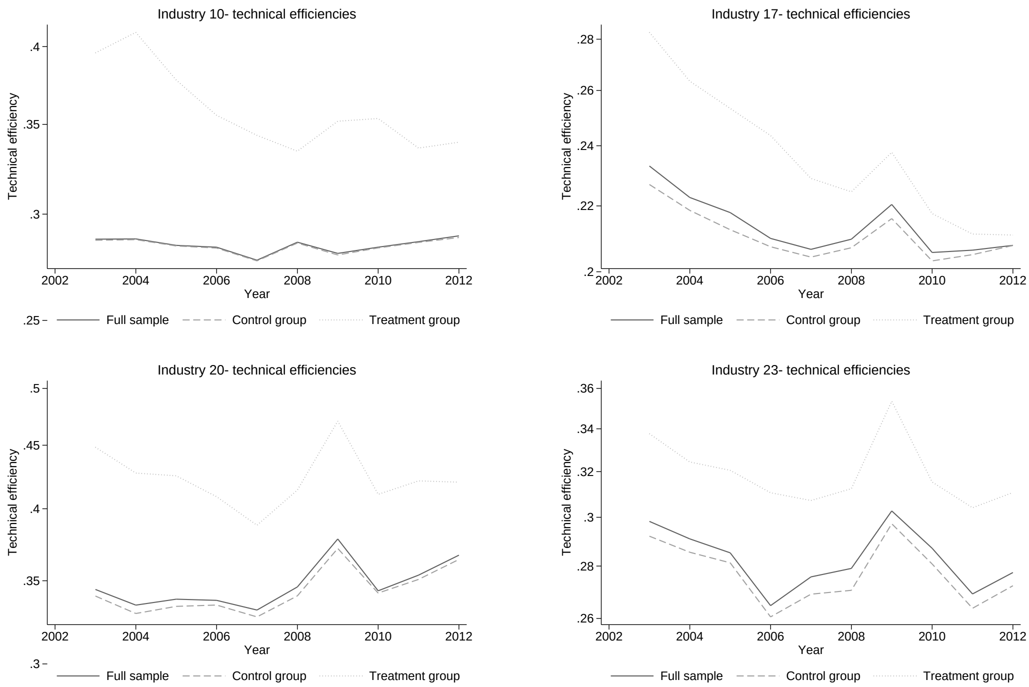

Notes: The vertical axis is displayed in log scale. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

3.5.2 Empirical evidence on identifying assumptions

In this section, we assess the validity of our key identifying assumptions that are described in Section 3.3 and derive strategies for our main analysis in order to take potential issues into account.

Common support: Table 3.4 shows descriptive statistics of the outcome variable technical efficiency and the observable covariates for the pretreatment year 2003. The upper panel refers to the treated firms, i.e. the firms regulated by the EU ETS, whereas the lower panel refers to the control firms. A comparison of the percentiles across groups indicates, that the common support assumption is reasonable for the outcome variable technical efficiency. The same comparison for the observable covariates reflects, however, the structural differences between treated and control firms. These differences occur due to the design of the EU ETS, that only regulates large emitters of CO_2.

In order to check the robustness of our difference-in-differences approach with regard to the common support assumption, we estimate a specification that includes firm-level fixed effects. This specification primarily relies on the within variation and thus is less prone to violations of the common support assumption. In the framework of our nearest neighbor matching approach, we take this issue into account by only considering observations that fulfill the common support assumption.

Unconfoundedness: For our baseline difference-in-differences specification, we assume that the counterfactual technical efficiency is equally distributed across treatment and control group conditional on the group, industry, and year fixed effects as well as interaction terms. We relax this assumption gradually by including additional observable firm characteristics and then firm-level fixed effects. We are able to investigate the validity of this assumption by analyzing differences in pretreatment trends of the outcome variable across groups. In particular, we apply our identification strategies to the pretreatment years assuming that the EU ETS was already introduced in 2004. The upper panel in Table 3.5 shows the resulting placebo treatment effects for our baseline difference-in-differences specification (Specification A) and the difference-in-differences specification including observable firm characteristics as covariates (Specification B). While our assumption of parallel trends in the absence of treatment holds for our subsample analysis, we see that for the full sample, there might be differences across treatment and control group

Table 3.4: Comparison of treated and control firms in 2003

	Mean	SD	P5	P50	P95	N
ETS firms						
Technical efficiency	0.416	0.324	0.177	0.335	0.861	473
Output (in EUR 1,000)	502,169.00	2,175,880.00	2,967.29	85,086.38	1,598,206.00	476
Emissions (in t CO_2)	266,627.70	845,970.20	2,697.63	53,453.96	1,018,856.00	475
Capital stock (in EUR 1,000)	223,274.10	785,493.00	1,548.98	40,658.98	807,826.10	477
Number of employees	1,844.19	7,239.40	28.83	352.33	7,512.33	477
Energy use (in MWh)	859,570.40	2,757,099.00	8,246.018	160,961.90	3,231,146.00	475
Non-ETS firms						
Technical efficiency	0.342	0.322	0.145	0.272	0.725	35,122
Output (in EUR 1,000)	22,675.17	262,125.30	912.28	5,187.19	76,554.40	37,412
Emissions (in t CO_2)	3,970.36	65,167.58	39.81	363.49	9,820.83	36,510
Capital stock (in EUR 1,000)	8,561.55	77,898.59	144.87	1,797.59	28,903.38	36,622
Number of employees	131.75	998.39	19.50	49.50	421.25	37,842
Energy Use (in MWh)	10,503.44	188,057.90	93.48	887.83	25,274.52	36,474

Notes: Output (production value), and capital stock are denoted in EUR 1,000. Energy use is denoted in MWh and CO_2 emissions in t CO_2. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

that are not captured by the fixed effects and the observational covariates. The placebo effect is economically small, but statistically significant. When interpreting the results of Specification B for the full sample, we have to take this into account. Furthermore, we add a specification with firm-level fixed effects to better control for differences across groups.

The lower panel in Table 3.5 shows the placebo treatment effects based on our matching approach for different numbers of nearest neighbors chosen by the Mahalanobis distance. None of the estimates is statistically significant indicating that the our conditional unconfoundedness assumption holds for the matching approach.

SUTVA: Our identification strategy relies on the assumption of stability of unit treatment values. It requires that the regulation by the EU ETS only affects regulated firms excluding spillover and equilibrium effects. This assumption cannot be directly tested. However, it is possible to estimate alternative specifications taking potential equilibrium effects into account. For our examination of the effects of the EU ETS, we differentiate between two cases: SUTVA could be either violated by equilibrium effects across or within industries.

Table 3.5: Pretreatment analysis

Parametric difference-in-differences model			
		Specification A	Specification B
Manufacturing		-0.0186 (0.0098)	-0.0095* (0.0038)
Food products (10)		0.0034 (0.0178)	-0.0057 (0.0072)
Paper and paper products (17)		-0.0178 (0.0275)	-0.0058 (0.0034)
Chemicals and chemical products (20)		-0.0376 (0.0259)	-0.0032 (0.0034)
Other non-metallic mineral products (23)		-0.0090 (0.0193)	-0.0018 (0.0031)
Nearest neighbor matching difference-in-differences			
	one neighbor	five neighbors	twenty neighbors
Manufacturing	0.0168 (0.0168)	0.0010 (0.0116)	0.0012 (0.0090)

Notes: Standard errors are computed by using the block bootstrap algorithm with 500 replications - robust with regard to heteroscedasticity and intra-firm correlation. * significant at the 5 percent level. A denotes the baseline specification, B includes explanatory variables. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

For the first case, consider for example a situation, where the EU ETS creates incentives for regulated firms to invest in abatement technology or new, more efficient machinery. As a consequence, the EU ETS does not only affect treated firms, for example in the cement or glass industry, but also indirectly potential control firms in other industries, such as manufacturing machinery and equipment. A similar line of thought is applicable to unregulated firms in the coking and refining industries, if regulated firms switch from carbon intensive to less carbon intensive fuels or energy (e.g. renewable energy sources). We aim to overcome this violation of SUTVA by examining the effect of the EU ETS within subsamples, in particular two-digit industries.

For the second case, the SUTVA violation within industries, consider for example a situation, where production is shifted from regulated to unregulated facilities. Fowlie, Holland, and Mansur (2012) use spatial variation in stringency of regulation. For our application, this is regrettably not feasible, since the EU ETS is uniformly applied to the regulated firms. We will discuss potential consequences of this kind of SUTVA violation for our results in Section 3.6

3.5.3 Difference-in-differences

The estimated treatment effects based on our three parametric difference-in-differences specifications are reported in Table 3.6. Specification A, B, and C refer to the baseline specification described in Equation 3.3, the specification including explanatory variables described in Equation 3.4, and the specification including explanatory variables and firm-level fixed effects described in Equation 3.5, respectively. All specifications include two-digit industry fixed effects, year fixed effects and complete interaction terms. The technical efficiency is computed as difference between the output predicted by the stochastic production frontier and the actual output of the firm. This distance to the frontier is positive for all firms. We take the natural logarithm of outcome variable and explanatory variables, the estimated treatment effects thus can be interpreted as semi-elasticities.

Table 3.6: Difference-in-differences treatment effects

	Specification A		Specification B		Specification C	
	03-07	03-12	03-07	03-12	03-07	03-12
Manufacturing	-0.0382*	-0.0510*	0.0003	-0.0003	-0.0052	-0.0042
	(0.0102)	(0.0130)	(0.0054)	(0.0066)	(0.0029)	(0.0039)
Food products (10)	-0.0284	-0.0667	-0.0058	-0.0066	-0.0036	0.0003
	(0.0274)	(0.0353)	(0.0039)	(0.0057)	(0.0035)	(0.0056)
Paper and paper products (17)	-0.0137	-0.0872*	-0.0139*	-0.0210*	-0.0134*	-0.0167*
	(0.0252)	(0.0299)	(0.0038)	(0.0047)	(0.0039)	(0.0044)
Chemicals and chemical products (20)	-0.0176	-0.0440	-0.0010	0.0009	-0.0074	-0.0048
	(0.0309)	(0.0426)	(0.0049)	(0.0057)	(0.0041)	(0.0056)
Other non-metallic mineral products (23)	0.0009	-0.0185	-0.0036	-0.0045	-0.0043	-0.0040
	(0.0161)	(0.0210)	(0.0025)	(0.0029)	(0.0024)	(0.0026)

Notes: Standard errors are computed by using the block bootstrap algorithm with 500 replications - robust with regard to heteroscedasticity and intra-firm correlation. * significant at the 5 percent level. A denotes the baseline specification, B includes explanatory variables, and C includes explanatory variables and firm-level fixed effects. All specifications include industry and time fixed effects and the full set of interaction terms. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

The first row of Table 3.6 shows the results for the entire manufacturing sector. The estimated treatment effects of Specification A indicate an economically and statistically significant negative impact of the EU ETS on technical efficiency. However, when we include addtional observable explanatory variables (Specification B) and firm-level fixed effects (Specification C), then the effect diminishes and becomes statistically insignificant. These results suggest, that the estimated treatment effects based on Specification A are biased due to confounding factors, which we are able to control for in the Specifications B and C.

Heterogeneity across industries within the manufacturing sector, for example with regard to abatement options and free allocation, might lead to insignificant treatment effects for the manufacturing sector as a whole. We therefore examine the estimated treatment effects for two-digit industries with sufficient observations in the treatment group. For the industries manufacture of food products, chemicals and chemical products, and other non-metallic mineral products (cement, glass, etc.), we do not find a significant effect of the EU ETS on technical efficiency of treated firms. Similar to the results based on the entire manufacturing sample, the estimates diminish, when controlling for confounding factors.

For the paper industry, however, we find statistically and economically significant treatment effects for all specifications and time periods considered. The size of the estimated treatment effects also decreases when controlling for confounding factors. Our preferred difference-in-differences model is Specification C indicating a -1.34 percent decrease in technical efficiency of regulated firms due to the EU ETS when considering only data until the end of Phase I and a -1.67 percent decrease when considering the data for both trading periods. In order to further investigate the better performance of EU ETS regulated firms in the paper industry, we show in Figure 3.3 the development of the indexed median of the output and inputs for treatment and control group separately. Figure 3.3 indicates, that the output of the treatment group increased more strongly in comparison to the control group during Phase I and Phase II. Furthermore, the treatment group conducted higher investments during the years 2007 and 2009 leading to a slightly higher capital stock during Phase II in comparison to the control group. The firms of the treatment group decreased employment after the investments in new capital stock. The energy use follows a similar trend across groups. This investigation of the descriptive statistics suggests that the difference in technical efficiency across groups mostly evolved due to increased output and investments in capital that is more efficient with regard to the use of employment.

3.5.4 Nearest neighbor matching

Table 3.7 shows the result of our nearest neighbor matching approach. Following Fowlie, Holland, and Mansur (2012), we implement a combination of nearest neighbor matching and difference-in-differences. Instead of parametrically accounting for observable con-

Figure 3.3: Comparison of treatment and control groups - the paper industry

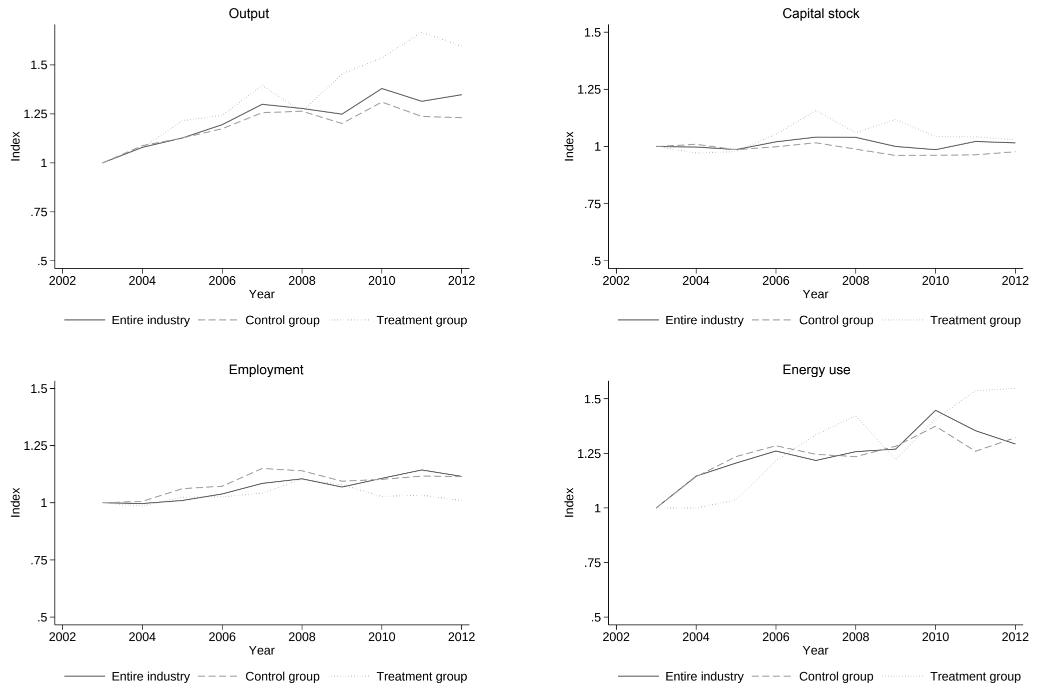

Notes: Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

founding factors, we identify an adequate control group using the Mahalanobis distance that determines similarity between firms by a weighted function of observable covariates for each firm. The weight is based on the inverse of the covariates' variance-covariance-matrix. This approach is nonparametric and does not assume a functional form for the outcome- or the treatment-model. The intuition behind this approach is to form a control group using unregulated firms that resemble the firms in the treatment group and thus might be affected by unobservable confounding factors in the same way. In line with Fowlie, Holland, and Mansur (2012), we apply nearest neighbor matching with replacement, i.e. unregulated firms can be used multiple times as a match.

We match on the firms' output, emissions, deployed capital stock, number of employees, and energy use in 2003 and match exactly on two-digit industries.[18] Table 3.7 shows estimated treatment effects for matching with the nearest neighbor, the five nearest neighbors, and the 20 nearest neighbors, respectively. The results should be interpreted jointly,

[18] Two-digit industries without treated firms are not considered in the estimation.

since a higher number of matched control firms improves the efficiency of the estimate, but at the same time introduces potential bias (Smith, 1997). The upper panel in Table 3.7 shows estimated treatment effects for year by year comparisons (base year is 2003). The lower panel shows estimated treatment effects for Phase I and Phase II. Apart from 2012, the estimated treatment effects are mostly negative. Only the year 2005 shows statistically significant effects that range between -2.77 and -3.43 percent. Pooling the data for the compliance periods, we find a significant negative effect of the EU ETS on firm specific technical efficiency during Phase I. The parameter estimates for the treatment effect in Phase II are of the same magnitude but statistically insignificant.

Table 3.7: Nearest neighbor matching treatment effects

	one neighbor	five neighbors	twenty neighbors
Year by year comparison (base year 2003)			
2005	-0.0158	-0.0343*	-0.0277*
	(0.0149)	(0.0134)	(0.0120)
2006	-0.0169	-0.0121	-0.0087
	(0.0153)	(0.0133)	(0.0134)
2007	-0.0152	0.0015	-0.0080
	(0.0224)	(0.0167)	(0.0154)
2008	-0.0171	-0.0001	0.0013
	(0.0247)	(0.0172)	(0.0288)
2009	0.0013	0.0069	-0.0018
	(0.0288)	(0.0222)	(0.0193)
2010	-0.0226	-0.0038	-0.0066
	(0.0293)	(0.0228)	(0.0200)
2011	-0.0021	-0.0129	-0.0082
	(0.0318)	(0.0295)	(0.0202)
2012	0.0190	0.0029	0.0122
	(0.0316)	(0.0334)	(0.0215)
Comparison trading periods with pretreatment period			
Phase I	-0.0289*	-0.0280*	-0.0265*
	(0.0124)	(0.0119)	(0.0108)
Phase II	-0.0294	-0.0097	-0.0164
	(0.0222)	(0.0181)	(0.0166)

Notes: Standard errors are robust with regard to heteroscedasticity and intra-firm correlation. * significant at the 5 percent level. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

3.6 Concluding Discussion

In this study, we investigate the effect of the EU ETS on the economic performance of regulated German manufacturing firms. We estimate a stochastic production frontier to recover the technical efficiency as firm specific measure for economic performance. Combining the difference-in-differences framework with parametric conditioning strategies and nonparametric nearest neighbor matching, we isolate the effect of the EU ETS on technical efficiency.

The results of the parametric difference-in-differences approach suggest that the EU ETS does not homogeneously affect firms in the manufacturing sector. We do not find a statistically significant effect of the EU ETS using data for the entire manufacturing sector. A subsample analysis at the two-digit industry level, however, shows that the EU ETS has a stronger influence on firms in particular industries. The industries manufacture of food products (10), manufacture of paper and paper products (17), manufacture of chemicals and chemical products (20), and manufacture of non-metallic mineral products (23) contain a sufficiently high number of regulated firms and enable us to examine the effect of the EU ETS on firms within narrowly defined industries. While we do not find a statistically significant effect of the EU ETS on the industries 10, 20, and 23, we find that the EU ETS significantly increased the economic performance of regulated firms in the paper industry.

The results based on the nonparametric nearest neighbor matching suggest a statistically significant positive effect of the EU ETS on the economic performance of the regulated firms during the Phase I of the EU ETS. A year-by-year analysis shows that the effect was only significant during the first year of Phase I. The EU ETS therefore had a particular strong effect when it was introduced.

Our results are in line with the results of the studies investigating the effect of the EU ETS on firms from the German manufacturing sector. Petrick and Wagner (2014) and Lutz (2016) do not find a statistically negative significant effect of the EU ETS on output, input use, and productivity. In contrast, Petrick and Wagner (2014) find a positive effect of the EU ETS on output while the inputs remain unaffected and Lutz (2016) finds a positive effect on productivity.

The results could be consistent with the following line of thought. The EU ETS might have incentivized investments in more efficient capital stock that allowed the firms to produce more output with less inputs. Alternatively, firms might have profited from free allocation and might have used the free resources to invest in more efficient capital stock.

When interpreting the results of our empirical analysis, it is important to bear in mind that we assume the EU ETS only to influence the treated firms. However, through spillover and equilibrium, effects the EU ETS might also have an impact on the economic performance of untreated firms. Conducting a subsample analysis, we can take into account equilibrium effects across industries, but we are not able to control for equilibrium effects within industries.

Furthermore, the design of our empirical strategy focuses on the identification of the EU ETS. We do not consider other regulatory instruments, such as energy taxes, that might interact with the effects of the EU ETS.

In order to overcome these caveats, it would be necessary to choose a different empirical strategy with additional assumptions on the underlying economic structure. This endeavor is left for future research. In addition, it would be interesting to apply our empirical strategy to production census data from other countries in order to assess the generality of our results.

3.7 Acknowledgements

We thank Ulrich Wagner, Kathrine von Graevenitz, Katrin Sommerfeld, and Goytom Abraha Kahsay for suggestions and insightful comments. We thank the participants of the Münster Energy Conference 2016 as well as participants of the seminars at the University of Münster and at ZEW in Mannheim. We gratefully acknowledge the Research Data Centre (FDZ) of the Federal Statistical Office and the Statistical Offices of the German Länder for granting us access to the AFiD data and for the use of their research facilities, in particular Michael Rößner for his advice and technical support regarding the estimation of our empirical model. Andreas Löschel, Benjamin Johannes Lutz, and Shunsuke Managi thank the Helmholtz Association for the financial support through the Helmholtz Alliance ENERGY-TRANS. The views expressed in this paper are those of the authors and do not necessarily represent those of the ZEW or the FDZ.

3.8 Appendices

Appendix 3.A. Data description

Industry classification: The underlying industry classification NACE rev. 2 (Statistical Classification of Economic Activities in the European Community) is the European implementation of the UN classification ISIC rev. 4. From 2003 to 2008 the data set contains the industry classification based on NACE rev. 1.1. For these years, we use the four digit industry codes and the official reclassification guide of the statistical offices (Quelle) in order to transfer NACE rev. 1.1 code to NACE rev 2.

CO_2 emissions: The Official Firm Data for Germany (Amtliche Firmendaten für Deutschland - AFiD) is a highly detailed data source with regard to energy use. The Energy Use Module contains information on the purchase, storage, sale, and use of 33 different fuels. We have access to slightly aggregated version of the Energy Use Module that contains information on 9 different fuels: natural gas, light fuel oil and heating oil, district heat, liquid gas, coal products, other mineral oil products, other gases, biomass, and other fuels. The Energy Use Module further includes information on the purchase, generation, sale, and use of electricity.

Following Petrick, Rehdanz, and Wagner (2011), we combine the energy use data from AFiD with data on CO_2 content in fuels and electricity. Table 3.8 shows the emissions coefficients we use in order to compute plant and firm-level CO_2 emissions. The coefficients for natural gas, light fuel oil, and liquid gas are directly taken from the official statistics of the Federal Environmental Agency (2012). The Federal Environmental Agency (2008) computes CO_2 emission coefficients for Germany in the years 2000 and 2005. We use the average coefficient over the two years. For the categories coal products, mineral oil products and other gases, we compute annual weighted averages in order to approximate adequate coefficients. The weights are determined by the sectoral use of the different fuels in the respective category and year (AG Energie Bilanz e.V., 2014). Our source for the electricity CO_2 coefficients is the official report Federal Environmental Agency (2014).

Table 3.8: CO_2 content of electricity and fuel use (g CO_2/kwh)

	03 - 12	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012
Natural gas	201.6										
Light fuel oil	266.4										
District heat	219.5										
Liquid gas	230.4										
Coal products		362.1	362.2	359.9	359.6	358.7	357.4	360.0	358.7	355.7	355.6
Mineral oil products		279.2	278.8	278.6	278.9	279.5	278.8	278.1	276.9	276.3	275.8
Other gases		195.9	195.9	195.9	195.9	195.8	195.9	195.9	195.8	195.5	195.6
Electricity		629	608	605	609	623	588	573	559	564	586

Notes: Sources: Federal Environmental Agency (2008), Federal Environmental Agency (2012), Federal Environmental Agency (2014), and AG Energie Bilanz e.V. (2014) , own calculations.

Appendix 3.B. Descriptive statistics

In Appendix 3.B, we show additional descriptive statistics. Figure 3.4 displays the price time series of the EUA futures traded at ICE. Figure 3.5 sheds some light onto the development of the firm characteristics over time. Each plot shows indexed medians for the according two-digit industry.

Figure 3.4: Price development of EUA futures

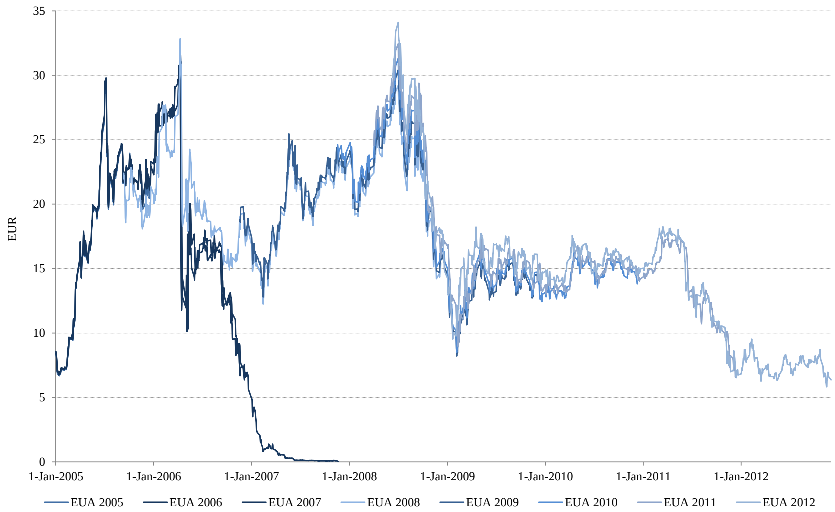

Notes: Source: ICE – accessed via Thomson Reuters Datastream.

Figure 3.5: Indexed medians for two-digit industries (I/II)

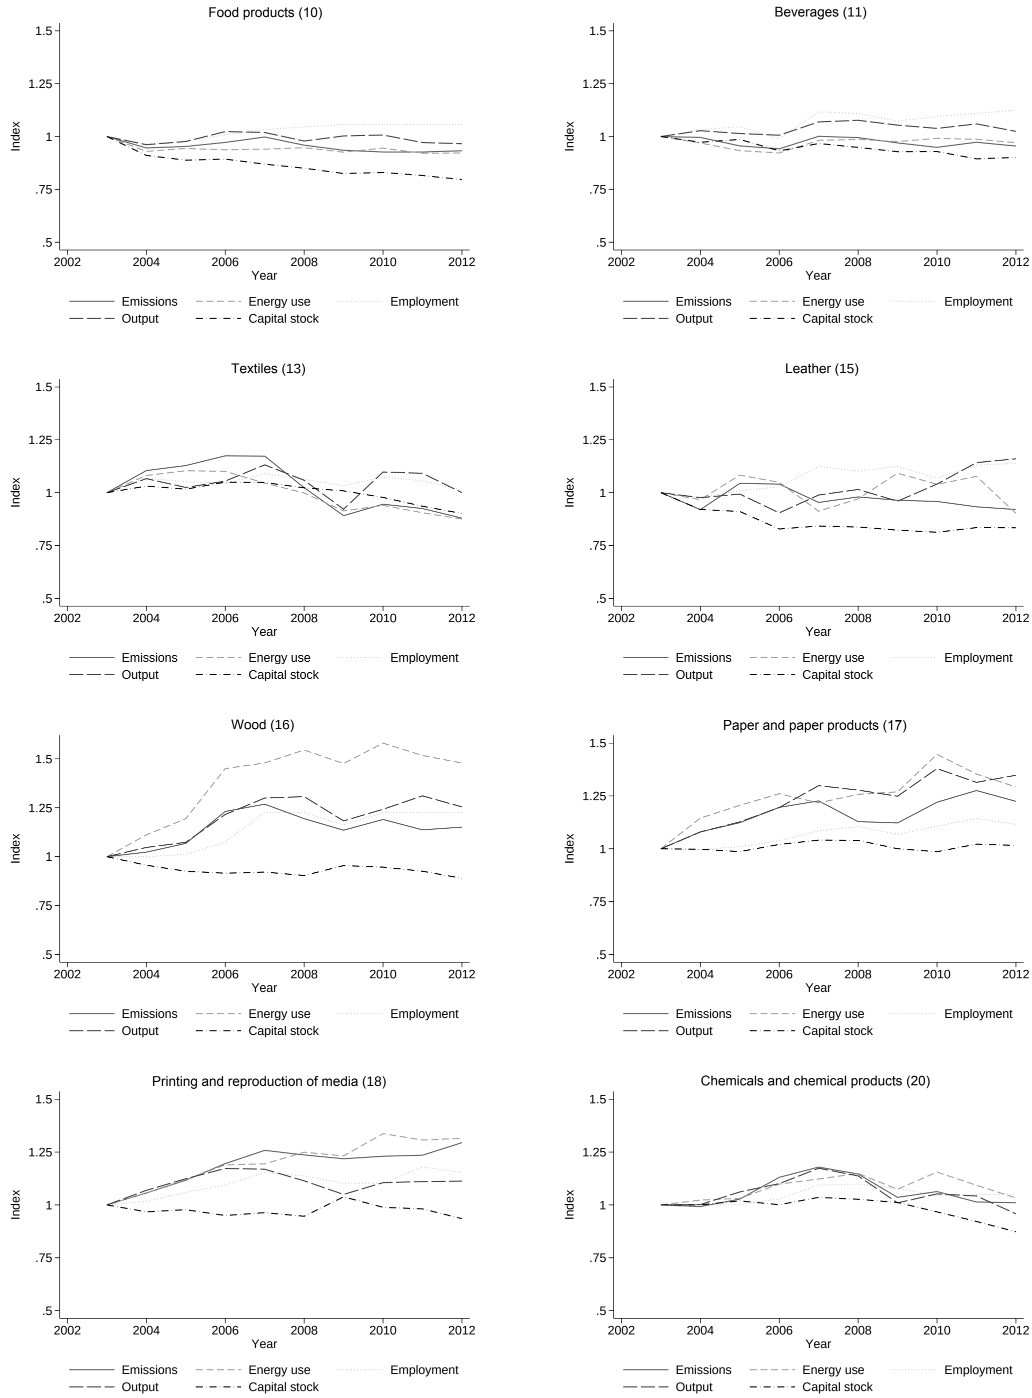

Notes: Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

Figure 3.6: Indexed medians for two-digit industries (II/II)

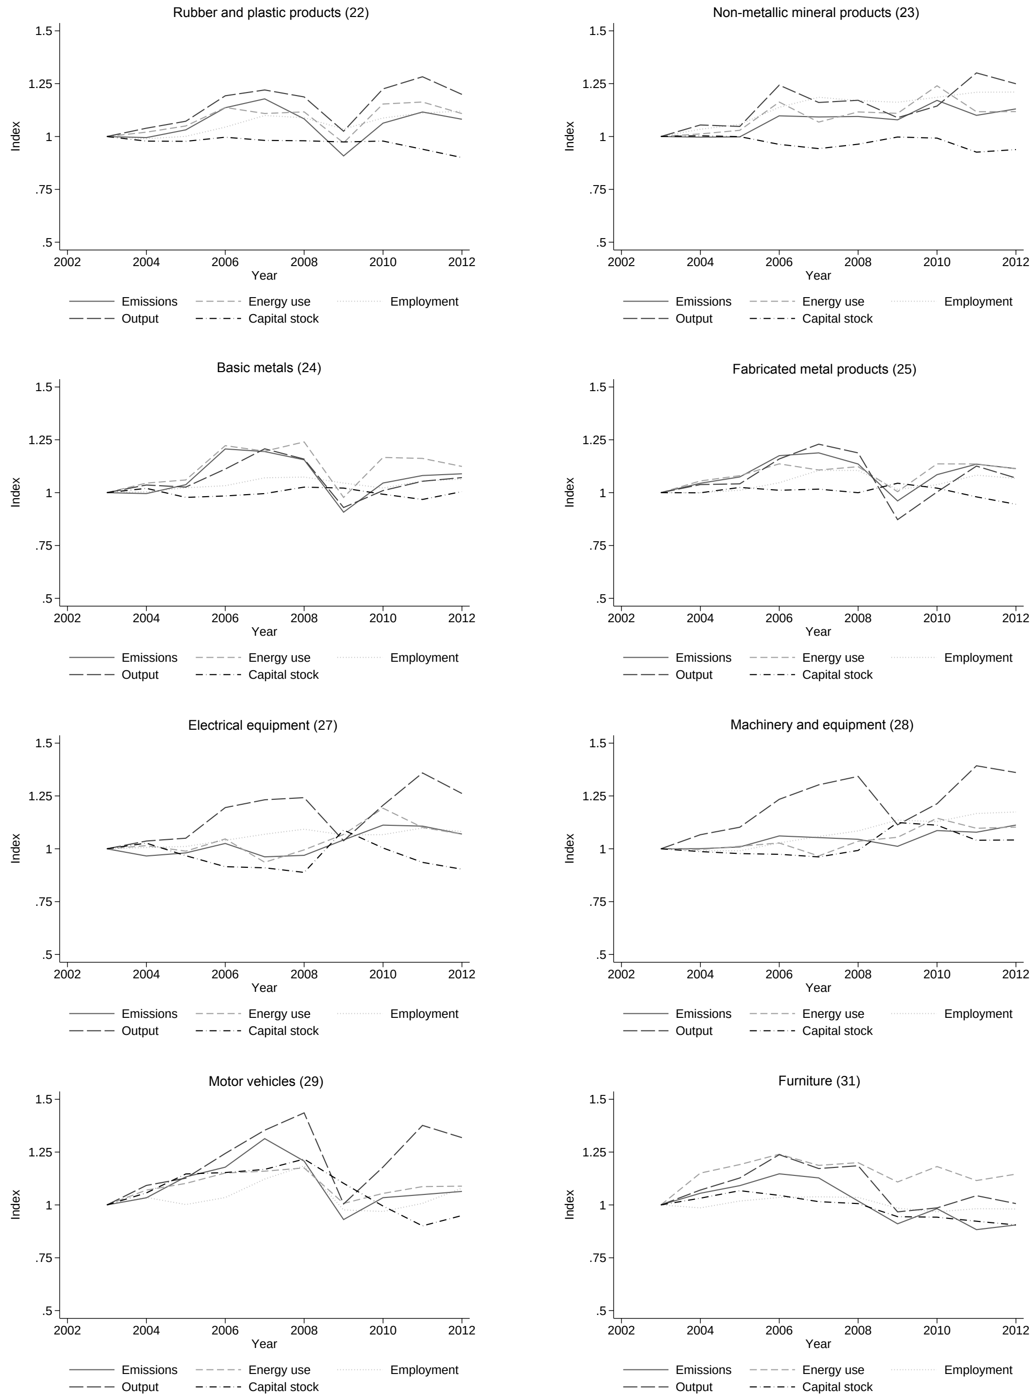

Notes: Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

Appendix 3.C. Recovery of technical efficiencies

In this appendix, we give a short description of the stochastic production frontier approach used to recover the technical effciencies.

According to Aigner, Lovell, and Schmidt (1977), the production process is influenced by a composite error term that consists of two economically distinguishable unknown random variables. The first component of the error term characterizes deviations from the optimal production frontier that result from decisions by the firm, e.g. mismanagement or suboptimal use of inputs. This random variable can be interpreted as a non-positive indicator for inefficiency. The second component of the error term captures noise and takes into account random factors that are not controlled by the firm, such as weather or unpredicted changes in the performance of machinery and employees, for example due to malfunction or illness. We follow Aigner, Lovell, and Schmidt (1977) and estimate the stochastic production frontier

$$\ln y_{it} = \ln f(\mathbf{x}_{it}) + \nu_{it} + u_{it}, \tag{3.6}$$

where y_{it} denotes the output of firm i at year t, $f(\mathbf{x}_{it})$ is the deterministic production frontier, $\mathbf{x}_{it}$ is a vector of inputs, ν_{it} is a nonpositive random variable depicting inefficiency, and u_{it} is a indpendently and identically distributed error term with zero mean and constant variance. We assume the deterministic frontier $f(\mathbf{x}_{it})$ to take the form of a Cobb-Douglas function. The vector of inputs $\mathbf{x}_{it}$ includes capital stock, labor and energy use. We assume the efficiency component ν_{it} to be drawn from a truncated normal distribution $\mathrm{N}^{+}(\mu_{\nu}, \sigma_{\nu}^{2})$ and the noise component u_{it} to be drawn from a symmetric normal distribution $\mathrm{N}(0, \sigma_{u}^{2})$. We implement the model using maximum likelihood estimation. In order to account for industry specific technologies, we estimate the stochastic frontier model for each two-digit industry within the manufacturing sector. The implied technical efficiency refers to a joint frontier for the years from 2003 to 2012.

Appendix 3.D. Alternative approach: value added stochastic production frontier

As a robustness check, we show results of an alternative stochastic production frontier. Instead of estimating a gross output production frontier, here, we estimate a value added production frontier. In addition, we show the results of the subsequent difference-in-differences analysis.

Table 3.9: Stochastic production frontier: value added specification

NACE	Industry	# Firms	Capital	Labor	Constant	σ_u	μ_ν	σ_ν
10	Food products	3687	0.313 (0.010)	0.661 (0.016)	3.052 (0.051)	0.513 (0.013)	-338.702 (4.423)	9.163 (0.315)
11	Beverages	536	0.162 (0.039)	0.936 (0.058)	3.255 (0.209)	0.715 (0.050)	-3.463 (0.120)	0.093 (0.003)
13	Textiles	905	0.171 (0.019)	0.903 (0.029)	3.067 (0.088)	0.366 (0.018)	-388.038 (81.580)	10.431 (1.388)
15	Leather and related products	205	0.241 (0.032)	0.851 (0.055)	2.704 (0.145)	0.417 (0.037)	-229.526 (13.599)	6.167 (1.001)
16	Wood and products of wood and cork	987	0.164 (0.016)	0.894 (0.024)	3.263 (0.063)	0.341 (0.013)	-367.242 (8.899)	9.846 (0.397)
17	Paper and paper products	878	0.223 (0.018)	0.844 (0.029)	3.061 (0.075)	0.338 (0.015)	-352.371 (13.586)	9.547 (0.566)
18	Printing and reproduction of recorded media	978	0.112 (0.021)	0.947 (0.035)	3.441 (0.084)	0.401 (0.034)	-266.503 (5.343)	7.135 (0.825)
20	Chemicals and chemical products	1494	0.259 (0.024)	0.802 (0.032)	3.320 (0.087)	0.444 (0.015)	-401.639 (10.496)	10.788 (0.410)
22	Rubber and plastic products	2228	0.187 (0.012)	0.876 (0.016)	3.143 (0.050)	0.339 (0.009)	-284.896 (13.465)	7.695 (0.379)
23	Other non-metallic mineral products	1601	0.237 (0.011)	0.795 (0.016)	3.239 (0.060)	0.361 (0.012)	-388.428 (15.707)	10.507 (0.486)
24	Basic metals	1098	0.187 (0.018)	0.865 (0.025)	3.488 (0.066)	0.370 (0.011)	-386.221 (8.959)	10.393 (0.368)
25	Fabricated metal products	4934	0.125 (0.007)	0.958 (0.010)	3.340 (0.031)	0.352 (0.007)	-264.760 (18.406)	7.150 (0.388)
27	Electrical equipment	2294	0.162 (0.016	0.903 (0.023)	3.388 (0.047)	0.379 (0.019)	-318.548 (7.684)	8.576 (0.405)
28	Machinery and equipment n.e.c.	5821	0.083 (0.006)	1.009 (0.009)	3.614 (0.027)	0.353 (0.007)	-317.811 (4.628)	8.588 (0.204)
29	Motor vehicles, trailers, and semi-trailers	1401	0.157 (0.015)	0.908 (0.021)	3.374 (0.052)	0.439 (0.021)	-335.859 (4.972)	9.031 (0.377)
31	Furniture	940	0.140 (0.014)	0.943 (0.020)	3.2079 (0.064)	0.328 (0.018)	-304.466 (6.713)	8.198 (0.436)

Notes: Number of firms for the first year of Phase I of the EU ETS (2005), the first year of Phase II (2008) and the last year of Phase II (2012) that is also the last year we observe. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

Table 3.10: Pretreatment: value added specification

	Specification A	Specification B
Manufacturing	-0.0186 (0.0098)	-0.0095* (0.0038)
Food products (10)	0.0034 (0.0178)	-0.0057 (0.0072)
Paper and paper products (17)	-0.0178 (0.0275)	-0.0058 (0.0034)
Chemicals and chemical products (20)	-0.0376 (0.0259)	-0.0032 (0.0034)
Other non-metallic mineral products (23)	-0.0090 (0.0193)	-0.0018 (0.0031)

Notes: Standard errors are computed by using the block bootstrap algorithm with 500 replications - robust with regard to heteroscedasticity and intra-firm correlation. * significant at the 5 percent level. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

Table 3.11: Treatment effect: value added specification

	2004	2005 - 2007	2005 - 2012
Specification A	-0.0471* (0.0206)	-0.0464* (0.0155)	-0.0048 (0.0194)
Specification B	-0.0429* (0.0191)	-0.0306* (0.0151)	0.0099 (0.0171)
Specification C	-	-0.0334* (0.0155)	-0.0528* (0.0145)

Notes: Standard errors are computed by using the block bootstrap algorithm with 500 replications - robust with regard to heteroscedasticity and intra-firm correlation. * significant at the 5 percent level. A denotes the baseline specification, B includes explanatory variables, C includes firm-level fixed effects. All specifications include industry and time fixed effects and the full set of interaction terms. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

Chapter 4

The effect of electricity taxation on German manufacturing: a regression discontinuity approach

4.1 Introduction

Many countries recognize that the use of energy is associated with negative environmental externalities and apply market-based policy instruments for internalizing social costs. For industrial energy users, countries often provide exemptions or compensation in order to prevent potentially adverse effects on firms' competitiveness. Despite the widespread use of market-based policy instruments, causal empirical evidence on their effects on firms' competitiveness is still scarce.

Germany established an ad-quantum excise tax, a market-based instrument, on electricity use in 1999. In this paper, we evaluate the causal effects of this electricity tax on the economic performance of firms in the manufacturing sector. The government was concerned that the new electricity tax might harm the competitiveness of German manufacturing firms. Therefore, it provided relief to firms in the form of reduced marginal tax rates. Without an international harmonization of energy taxes, increasing electricity prices in Germany might have encouraged a relocation of electricity intensive production to countries with less stringent regulation. In fact, the Federation of German Industries has campaigned for low electricity prices over the last two decades: "Energy policy should not become a risk factor for Germany as a location for business and investment. Policy makers have to change the current course and must recognize electricity as an important factor of competitiveness." [1] While many German manufacturing firms are export-oriented and thus face strong international competition, energy costs are only a small share of the total costs for them, apart from few very energy intensive industries. Despite an ongoing public debate about potential competitiveness impacts of taxes on electricity use, the causal effects of the electricity tax on economic firm performance have not been evaluated so far.

The electricity tax varies in marginal tax rates allowing us to identify and estimate the causal effects of the electricity tax on the economic performance of German manufacturing firms. Using a sharp regression discontinuity design (Lee and Lemieux, 2010), we investigate how firms' turnover, exports, value added, investment, and employment responded to different marginal tax rates. The marginal electricity tax rates are a deterministic and discontinuous function of firms' electricity use. Firms that use more electricity than certain thresholds established by legislation pay reduced marginal tax rates. These reduced marginal rates generate local random experiments at the thresholds from which they apply. The sharp nonparametric regression discontinuity design exploits the quasi-random variation in marginal electricity tax rates around the thresholds and allows us to identify and estimate the causal effects of the differential tax rates. Thereby we can evaluate the effectiveness of the compensation scheme, i.e., the reduced marginal tax rates. The difference between marginal tax rates in some years is larger than the full tax rate in other years, so we can also infer the effect of the electricity tax itself.

We make use of official micro-data on the activities of the German manufacturing at the plant and firm level. The data is collected by the German Federal Statistical Office through a rigorous census of firms on production, costs, and energy use. Participation in the surveys is mandatory by law for all plants with more than 20 employees. The surveys include detailed information about electricity use at the plant and firm level. Given that the electricity tax law specifies the marginal tax rate to be a deterministic function of electricity use, this allows us to calculate for each firm the electricity tax rate that applies.

The results suggest that the effects of the electricity tax on firms' turnover, exports,

[1] The Federation of German Industries represents industrial firms in Germany. It communicates to political decision-makers and the public on behalf of its 36 sector associations, acting for over 100,000 firms with eight million employees. The cited statement can be found in the official publication Federal Association of German Industry (2008).

value added, investment, and employment are neither systematic nor statistically significant. Gradually shifting the thresholds from which reduced tax rates apply may increase revenues for the government without adversely affecting the economic performance of firms. The additional tax revenues could be used to lower taxes that are widely regarded as particularly harmful to economic efficiency and growth such as taxes and social security contributions on labor, to consolidate budgets, or to finance new investments.

Our study contributes to an emerging literature on the causal effects of market-based environmental policy regulation by examining the case of the German electricity tax. Despite widespread regulatory intervention, there are so far only few studies that investigate the causal impact of market-based environmental regulation on environmental and economic performance of manufacturing firms. Using a quasi-experimental research design with a generalized matching estimator, Fowlie, Holland, and Mansur (2012) examine the effectiveness of Southern California's NO_x trading program that has been introduced in the framework of the Clean Air Act Amendments of 1990. They show that the tradable permit system yielded emission reductions of 20 percent in comparison to the counterfactual, where facilities were regulated by command-and-control regulation. Martin, de Preux, and Wagner (2014) evaluate the impact of a carbon tax on the manufacturing industry in the UK using an instrumental variable approach. They provide robust evidence that the Climate Change Levy significantly decreased energy intensity and electricity use, while the economic performance of the firms remained unaffected. Petrick and Wagner (2014) investigate the effect of the EU Emissions Trading System (EU ETS) on German manufacturing firms with the help of semi-parametric matching estimators. They find that the scheme curbed the CO_2 emissions by improving energy efficiency and fuel switching. According to their results, the scheme had no impact on economic performance of the regulated firms. Wagner, Muûls, Martin, and Colmer (2014) investigate the economic and environmental impact of the EU ETS on French manufacturing plants. Their results suggest a significant negative causal effect of the EU ETS on emissions of 15-20 percent. With regard to its effect on economic outcomes, they find a significant reduction of employment by 7 percent in regulated plants.[2]

In the following, we first explain how the design of the German electricity tax leads to

[2]For a survey of the literature on the effectiveness and efficiency of pricing carbon and in particular the EU ETS see Arlinghaus (2015) and Martin, Muûls, and Wagner (2016).

variation in firms' marginal electricity tax rate. Second, we discuss how we can identify and estimate the effects of the German electricity tax using a regression discontinuity design. Third, we describe the official data used in our analysis, which is collected by the German Statistical Office. Fourth, we present the results of our analysis and examine the robustness of our findings. We briefly discuss the implications of our results before we conclude.

4.2 The German electricity tax and variation in the marginal tax rate

The German electricity tax was introduced in 1999 aiming at improving energy efficiency and lowering labor costs. The new electricity tax increased the price on electricity incentivizing firms to reduce electricity use. The revenues lower social security contributions uniformly across firms and thereby, overall labor costs. We aim to assess how differences in marginal electricity tax rates affected firms' economic performance.

The electricity tax is levied on electricity use as an ad-quantum excise duty with a full rate of 20.5 euros per MWh at present. This implies an effective tax on the carbon content in the average unit of electricity of 44.4 euros per tonne of carbon dioxide (CO_2). Although this calculation assumes that the generation mix of electricity would not change, if the tax was levied on CO_2 instead of on electricity, it indicates the significance of the electricity tax.

A comparison of the retail prices and the full rate shows that the tax significantly increases the retail price, between 27.1 percent in 2002 and 15.2 percent in 2005. Figure 4.1 shows the development of retail prices for electricity use and the full tax rate for the period from 1995 - 2005. The average price faced by a firm that consumes 2,000 MWh per annum ranged between 65 and 100 euros during this time period (Eurostat, 2014), which we take as the lower bound of the electricity price. As the upper bound of the electricity price, we show the price for a household that consumes 3.5 MWh per annum, which ranged between 115 and 135 euros (Eurostat, 2014).

The government was concerned that the electricity tax may harm the competitiveness of German firms that are subject to competition from abroad. For that reason, the government took at least two measures. First, it introduced the electricity tax in several steps

Figure 4.1: Retail prices for electricity 1995 - 2005

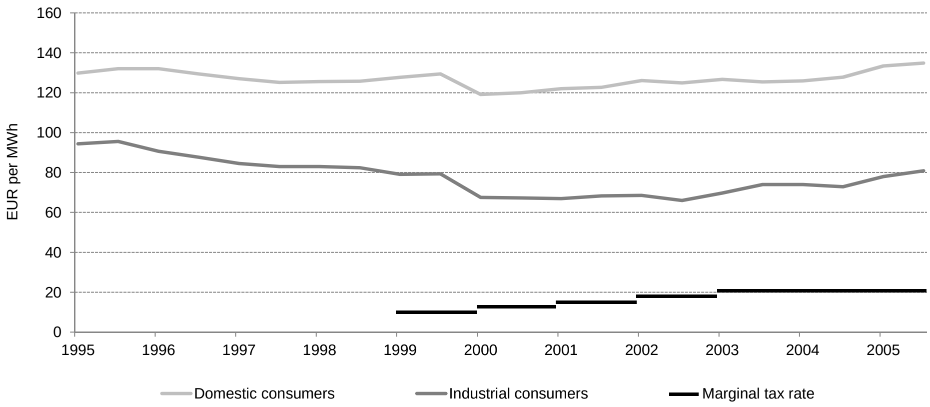

Notes: The price for domestic electricity use relates to a reference household that annually consumes 3.5 MWh of which 1.3 MWh are consumed at night. The price for industrial electricity use relates to a reference firm that annually consumes 2,000 MWh (max. demand 0.5 MW; annual load 4,000 hours). Prices are denoted in EUR per MWh, include transmission, system services, meter rental, distribution and other services and exclude taxes and levies. Source: Eurostat (2014), own calculations.

until the full rate was reached in 2003 giving firms time to adjust to higher electricity prices. Second, it provided relief to manufacturing sectors through reduced tax rates.

The reduced tax rates apply from certain thresholds of electricity use onwards and are key to our identification strategy as subsequently outlined and described more formally in Section 4.3.1. While every user has to pay the same marginal tax rate for any use below the threshold, firms in the manufacturing sector are eligible for a reduced marginal tax rate for any use above the threshold. Table 4.1 shows that the tax is a piecewise linear function of electricity use X, that can be characterized as a set of two linear taxes, each relevant to only a particular range of X. Let $t(0)$ stand for the regular marginal tax rate and $t(1)$ for the reduced marginal tax rate. The known threshold, from which the reduced marginal tax rate applies, is denoted by c. Then, the tax function can be written as

$$T(X) = \begin{cases} t(0)X & \text{if } X \leq c \\ t(1)(X-c) + t(0)c & \text{if } X > c \ . \end{cases} \tag{4.1}$$

The thresholds of 50 MWh or lower for a reduced marginal electricity tax rate may seem low; nevertheless many firms in the manufacturing sector consume about that much electricity. In 1999, when the electricity tax was introduced, about 25.2 percent of the firms in the data set used less than 100 MWh electricity per annum and about 13.1 percent of the firms used less than 50 MWh electricity per annum (see also Figure 4.2 in Section 4.4.3).

Thus, many firms in the manufacturing sector consume about that much electricity, and are therefore directly affected by either having to pay the reduced or the full marginal electricity tax rate.

The reduced marginal tax rate for any electricity use above the threshold in a given year generates random variation in firms' marginal electricity tax rates. Whether firms face the full or reduced marginal tax rate is essentially chance due to arbitrarily set thresholds. We use this random assignment to identify the effects of the reduced marginal tax rates on firms' economic performance with a regression discontinuity design as explained in the following section.

Another type of electricity tax reduction is the so-called Spitzenausgleich. Recall that the revenues from the electricity tax are used to lower social security contributions on labor uniformly across firms. While firms benefit from reduced social security contributions, they may eventually face overall higher costs due the new electricity tax. The Spitzenausgleich reimburses a certain percentage of the potential additional burden from the new electricity tax net of the savings on social security contributions. The reimbursement rule and also the reduction in social security contributions have changed several times.

Table 4.1: Marginal electricity tax rate

	Marginal electricity tax rate in EUR per MWh						
Electricity use threshold	Until 1999	1999	2000	2001	2002	2003	Until 2010
below 25 MWh	0	10	12.5	15	17.9	20.5	20.5
above 25 MWh	0	10	12.5	15	17.9	12.3	12.3
above 28.6 MWh	0	10	12.5	15	3.6	12.3	12.3
above 33 MWh	0	10	12.5	3	3.6	12.3	12.3
above 40 MWh	0	10	2.5	3	3.6	12.3	12.3
above 50 MWh	0	2	2.5	3	3.6	12.3	12.3

The Spitzenausgleich applies only for electricity use above the same thresholds from which the reduced marginal electricity tax rate is granted. Thereby it may add to the potential effects of the reduced marginal tax rates. We expect that the effects of the reduced tax rate dominate around the thresholds given non-negligible administrative procedures for receiving the Spitzenausgleich. In the following, we will therefore refer to the effects of the reduced tax rate, bearing in mind that some of effects may have been reinforced by the Spitzenausgleich.

In August 2006, exemptions to the electricity tax were granted for firms in the manufacturing sectors for the electricity consumed in various production processes. In particular, electricity used for electrolysis, production of glass, ceramics, fertilizers, metal production and processing, as well as chemical reduction processes was exempted from the electricity tax. The tax exemptions apply for all electricity consumed and thus not only from above certain threshold onwards. We do not have any information on how much electricity is used for these processes. From 2006 onwards, a clean identification of firms that benefit from the reduced marginal electricity tax rate is not possible any more. We therefore analyze the effects of the reduced marginal electricity tax rate only until 2005.

As mentioned, the revenues from the electricity tax are used to lower social security contributions. Given that the reduction of social security contributions applies to all firms uniformly, we cannot measure the effect of the reduction in social security contributions. Neither can we assess the overall effect of the reform package, i.e., the introduction of a new electricity tax combined with the use of its revenues to lower social security contributions. What we aim to assess is how different marginal electricity tax rates affected firms' economic performance.

4.3 Research design

4.3.1 Empirical approach

Our goal is to identify the causal effect of the electricity tax on the economic performance of firms in the manufacturing sector. As ad-quantum excise duty, the electricity tax increases the price per unit of consumed electricity by the marginal tax rate t. We build our identification strategy on variations in the marginal tax rate. Firms that are energy intensive in terms of individual electricity use face a lower marginal tax rate in comparison to less energy intensive firms. In particular, the reduced tax rate applies, if the electricity use X_i of firm i exceeds the known threshold c that is set by the regulatory authorities:

$$t_i = \begin{cases} t_i(0) & \text{if } X_i \leq c \\ t_i(1) & \text{if } X_i > c \ , \end{cases} \tag{4.2}$$

where $t_i(0)$ denotes the regular marginal tax rate and $t_i(1)$ the reduced marginal tax rate, respectively. Hence, the tax reduction scheme creates a sharp discontinuity in the marginal tax rate as a function of the individual electricity use. This feature of the electricity tax allows us to identify and estimate the effect of the electricity tax for any given year by employing a sharp regression discontinuity design.

The profit maximizing firm equalizes marginal costs and marginal revenues by choosing the level of output and the combination of inputs subject to technological constraints. The discontinuity in the marginal tax rate and the resulting scheme of two different marginal tax rates creates variation across firms regarding the marginal costs associated with the use of electricity. We expect that firms react to the regular and reduced marginal tax rate differently by adjusting the level of output and combination of inputs according to the marginal tax rate they face.

More specifically, we hypothesize that firms that face higher marginal taxes will have lower output relative to firms with low marginal costs. Two observations lead to this hypothesis. First, firms that have to pay the full tax rate face higher marginal costs for electricity use and thus, overall higher marginal costs than firms that only need to pay the reduced tax rate. For minimizing costs, a firm equates the ratio of marginal costs of inputs to the ratio of the marginal products of input factors. A higher marginal cost for electricity use translates into higher overall costs for producing the same level of output. Thereby, overall marginal costs are also higher for firms with higher marginal costs for electricity use. Second, if there are two types of firms in the market, those with low marginal costs are expected to produce a higher output than those with high marginal costs all else equal.

The economic outcomes we can observe with our dataset are firms' turnover, exports, value added, investment and employment. We expect that the turnover and exports of firms with the reduced tax rate will be higher than for those that face the full marginal tax rate. The intuition is that lower marginal costs allow the former firms to produce more. For the same reason, we also expect that the value added, which is revenue minus costs, of firms with the reduced tax rate is higher than for firms with the full marginal tax rate.

The total effects of the reduced marginal tax rate on investment and employment can have either sign. With regard to investment, there is a direct effect, namely that higher production causes more investment. Yet, there is also an indirect effect in the opposite direction. Firms that face high marginal costs due to paying the full tax rate have an incentive to invest in new, more energy-efficient production technology to mitigate their cost disadvantage. Thus, the total effects may have either sign. Regarding employment there is, first, a direct effect from lower marginal costs to higher production and thus, more employment. Second, there are indirect effects in addition, if firms with high marginal costs invest in new, more energy efficient technology. This new technology could either be less or more labor intensive than the old one. If it is less labor intensive, the indirect effect goes in the same direction as the direct effect and we thus expect firms with the reduced tax rate to employ more labor. If the technology is, however, more labor intensive than the old one, the indirect effect goes in the opposite direction, i.e. firms that pay the full tax rate employ more labor. In total, we cannot hypothesize unambiguously what the effect of reduced tax rate on labor is.

Our identification strategy can be formalized using the potential outcomes framework introduced by the seminal work of Rubin (1974, 1977). The firms of the German manufacturing industry are assigned to two different groups. The binary variable $D_i \in \{0, 1\}$ describes the treatment status of firm i. Let $D_i = 1$ if the firm' s electricity use X_i exceeds the threshold c. Then, the firm is subject to the reduced marginal tax rate $t_i(1)$ and is considered as treated. Let $D_i = 0$ if the firm's electricity use X_i is lower than the threshold c. In this case, the full marginal tax rate $t_i(0)$ applies and the firm is assigned to the control group. Consequently, we denote the potential outcomes by

$$Y_i = \begin{cases} Y_i(0) & \text{if } X_i \leq c \\ Y_i(1) & \text{if } X_i > c \ . \end{cases} \tag{4.3}$$

As shown in Equation 4.1, the assignment to the treatment group is a deterministic function of the electricity use X_i. Since we observe the electricity use X_i, we are able to identify if firm i belongs to the treatment or the control group. Following the sharp regression discontinuity design framework outlined by Imbens and Lemieux (2008) and Lee and Lemieux (2010), we analyze the sharp discontinuity in the conditional expectation of the outcome given electricity use X_i to unveil an average causal effect of the treatment:

$$\tau = \lim_{x \downarrow c} E[Y_i \mid X_i = x] - \lim_{x \uparrow c} E[Y_i \mid X_i = x]. \tag{4.4}$$

In the literature, this term is interpreted as the local average treatment effect at the threshold (Imbens and Lemieux, 2008):

$$\tau = E[Y_i(1) - Y_i(0) \mid X_i = c]. \tag{4.5}$$

Making use of assumptions we describe in Section 4.3.2, the treatment variation close to the threshold c is considered as good as random. The random assignment implies that the discontinuity at the arbitrarily set threshold c identifies the treatment effect of interest. Consequently, we are able to identify the effect of the electricity tax reduction by comparing firms of the treatment and control group that are in the neighborhood of the threshold.

We assume the annual electricity use of a firm to be mainly determined by its physical capital stock and the utilization of the installed capacities. The size of the capital stock and the deployed technologies are highly path dependent, i.e. the current state is the result of past investment decisions. The capital stock can only be changed through investment or disinvestment. Thus, it is fixed in the short run.

While the utilization of the installed capacities can be manipulated by the firm in the short run, utilization is also driven by stochastic shocks that cannot be controlled by the firm. These shocks might be of external origin, as for instance the effect of weather on heating or cooling processes, or of internal origin, as for instance the breakdown of a machine. The effect of these stochastic shocks might be small, but prevent firms from precisely manipulating electricity use.

We assume initially that firms make accurate production decisions based on the correct marginal costs associated with electricity use, and later relax this assumption. As long as the stochastic component in electricity use is fairly small, i.e. the error with regard to their estimate of the annual electricity use is small, a violation of this assumption is not severe for our identification strategy. Unfortunately, we do not have the possibility to empirically investigate whether firms are able to correctly predict their electricity use for the year ahead. In Section 4.5.7, we suggest an alternative model specification that relaxes the assumptions described here.

The tax reduction scheme is implemented through reimbursement, i.e. firms whose electricity use exceeds the threshold may request reimbursement from the local tax and custom agency. We do not observe whether firms that were assigned to the treatment group received the treatment. While the reimbursement procedure creates imperfect compliance, inference is still possible. We account for this case of *encouraged* treatment by performing an intent to treat analysis. We compare control and treatment group without regards to whether the tax reduction was actually claimed. Accordingly, the local average treatment effect measures in our case how the treatment *assignment* affected the firm's activities, as opposed to the desired measure of how the treatment *itself* affected the firm's activities (Pearl, 2000). For simplicity, we will stick with the term local average treatment effect. Yet, one should bear in mind that the estimated treatment effect measures the intend to treat, i.e. the effect of the eligibility for the electricity tax reduction.

4.3.2 Identifying assumptions

The regression discontinuity design allows us to identify local treatment effects under comparatively lax assumptions. Following Hahn, Todd and van der Klaauw (2001) and Lee and Lemieux (2010), we unfold the assumptions that underlie the approach and discuss them in light of the German electricity tax.

Assignment to the treatment group

First, the treatment assignment must be a monotone deterministic function of the assignment variable. This holds in our case, as firms that consume more electricity X_i than the known threshold c benefit from the tax reduction and are considered as treated, while firms that consume less face the full marginal tax rate (see Equation 4.2) and are considered as untreated. Second, the probability of treatment has to be a discontinuous function of the assignment variable. The probability to be treated, i.e. to benefit from the tax reduction, changes discontinuously at the threshold c, particularly $P[D_i = 1 \mid X_i = x]$ is 0 for $x \leq c$ and 1 for $x > c$.

Inability to precisely control the assignment variable

The central assumption that underlies our identification strategy is that firms cannot *precisely* manipulate their individual electricity use. Lee (2008) shows that the treatment in the regression discontinuity design is random, if the assignment variable has a continuously distributed stochastic component, i.e. firms cannot precisely control their electricity use. We argue that this assumption is plausible in our setting for two reasons: First, complex production processes and path dependencies in the manufacturing industry make *precise* manipulation of a firm's electricity use difficult. Second, exogenous factors that drive electricity use lead to stochastic variation in electricity use. For example, weather conditions or the breakdown of a machine might impact a firm's energy use. We will test this assumption in Section 4.5.1 by examining the empirical distribution of the assignment variable. No evidence for precisely controlling the assignment variable is found.

Local continuity restriction

In absence of treatment, the outcome variable has to evolve continuously with the assignment variable in the neighborhood of the threshold. If other factors create discontinuities in this relationship, a clear identification of the local treatment effect is not possible. In Section 4.5.1, we empirically investigate the evolution of each outcome variable as a function of the assignment variable electricity use for the years before the implementation of the electricity tax. In this way, we aim to detect other sources that create discontinuities in the relationships under investigation and thus might affect identification. No evidence for any prior discontinuities is found.

Stable unit treatment value assumption

The stable unit treatment value assumption (SUTVA) assumes, that the treatment status of a firm does not affect the outcomes for other firms. Hence, SUTVA excludes spill overs and general equilibrium effects across firms. This assumption cannot be directly tested. However, in Section 4.6, we will discuss the robustness of our results with regard to a potential violation of this assumption.

4.3.3 Estimation

The estimation of the local average treatment effect τ requires an estimator that shows good small sample properties and is suitable for inference at the boundary of the support of the regression function (here threshold c). Addressing these obstacles, Hahn, Todd, and van der Klaauw (2001) and Porter (2003) propose a nonparametric approach based on weighted local linear or polynomial regressions at both sides of the threshold. This estimator has become the standard choice for the estimation of local average treatment effects in the regression discontinuity literature. Yet, the estimator requires the selection of a bandwidth that determines the range around the threshold that is exploited for the estimation of the local regressions. We use a fully data-driven bandwidth algorithm developed by Imbens and Kalyanaraman (2012) in order to select the asymptotically optimal bandwidth.

We formalize the estimator of the local average treatment effect $\hat{\tau}$ at the threshold c

as described in Imbens and Kalyanaraman (2012):

$$\hat{\tau} = \hat{\alpha}_{+} - \hat{\alpha}_{-} \tag{4.6}$$

where $\hat{\alpha}_{+}$ and $\hat{\alpha}_{-}$ denote the constants of a weighted local linear regression. The weights are computed by applying a kernel function $K(\cdot)$ on the distance of each observation's score to the threshold c. The parameters are obtained by estimating two equations within two narrow windows left and right of the threshold that yield in the estimator $\hat{\alpha}_{+}$ for only treated and the estimator $\hat{\alpha}_{-}$ for only control firms:

$$(\hat{\alpha}_{+}, \hat{\beta}_{+}) = \underset{\alpha,\beta}{\operatorname{argmin}} \sum_{i=1}^{N} \mathbf{1}_{X_i>c}(Y_i - \alpha - \beta(X_i - c))K\left(\frac{X_i - c}{h}\right), \tag{4.7}$$

$$(\hat{\alpha}_{-}, \hat{\beta}_{-}) = \underset{\alpha,\beta}{\operatorname{argmin}} \sum_{i=1}^{N} \mathbf{1}_{X_i<c}(Y_i - \alpha - \beta(X_i - c))K\left(\frac{X_i - c}{h}\right), \tag{4.8}$$

where $\mathbf{1}_u$ is an indicator function taking the value 1 if condition u is fulfilled. In order to select the optimal bandwidth h for the two windows, we employ the algorithm developed by Imbens and Kalyanaraman (2012). The default form of the kernel function $K(\cdot)$ in our set up is triangular. The computed standard errors are robust with respect to heteroscedasticity and show good finite sample properties.[3] Unless otherwise stated, the results that are presented in the remainder of this paper are estimated based on the procedure shown in Imbens and Kalyanaraman (2012).

4.4 Data

4.4.1 Official Firm Data for Germany

Our empirical analysis exploits official census micro-data of firms collected by the German Federal Statistical Office and the Statistical Offices of the German Federal States. The data are confidential but the German statistical offices provide remote data access to researchers for scientific purposes. Participation in surveys conducted by the German statistical offices is mandatory by law and many official German government statistics build on this data.

[3] The estimator of the local average treatment effect shown here is implemented using the STATA package developed by Calonico et al. (2014a). For the computation of the standard errors, we choose the conventional fixed-matches variance estimator proposed in Calonico, Cattaneo, and Rocio (2014a, 2014b).

The dataset, called Amtliche Firmendaten für Deutschland - AFiD (Official Firm Data for Germany), records activities of the industrial sector at the plant and firm level. It consists of several modules, which can be combined. In particular, we use two modules that capture activities of the German manufacturing industry.

The core of our dataset is the module AFiD-Panel Industrial Units. This longitudinal census combines annual results from the Monthly Report on Plant Operation, the Census on Production, and the Census on Investment. The AFiD-Panel Industrial Units is a census of all establishments - physical buildings or structures, i.e., plants. It provides detailed information on turnover, exports, employment, investment, and firm affiliation.

This database is extended by the AFiD-Module Use of Energy. The AFiD-Module Use of Energy is a longitudinal census that comprises results from the Monthly Report on Plant Operation and the Census on Energy Use. It contains information about sale, purchase, generation, use, and distribution of electricity and fuels. Both the AFiD-Panel Industrial Units and the AFiD-Module Use of Energy have the same group of respondents: All German plants that operate in the manufacturing industry and belong to firms that employ more than 20 persons must participate in the census.

Merging the AFiD-Panel Industrial Units with the module AFiD-Module Use of Energy, we construct a data set comprising longitudinal census data at the firm level covering a time span from 1995 to 2005. This data cover pre-reform, reform, and post-reform periods. Where necessary, we aggregate plant-level data to the firm level using the firm affiliation provided within the AFiD-Panel Industrial Units.

4.4.2 The Cost Structure Survey

We link the AFiD-Panel Industrial Units and the AFiD-Module Use of Energy with data from the Cost Structure Survey (CSS) to obtain information on the value added at the firm level.[4]

The CSS also conducted by the German Federal Statistical Office and the Statistical Offices of the German Federal States gives detailed information on the costs from capital, labor as well as value added of firms on an annual basis from 1999 - 2005.

In contrast to the AFiD-Panel Industrial Units and the AFiD-Module Use of Energy,

[4]In particular, we use the variable gross value added - for practical purposes referred to as value added throughout the paper.

the CSS collects data directly at the firm level. It includes all firms with more than 500 employees. For firms with at least 20 and less than 500 employees, the statistical offices collect a random sample that is stratified by the number of employees and industry affiliation. These firms remain four years in the panel and are replaced by a new random sample afterwards. For the CSS, the same participation rules apply as for AFiD. The provision of the requested information is mandatory by law.

4.4.3 Descriptive statistics

In our analysis, we focus on German firms that belong to the sectors mining and quarrying (ISIC 1010-1429) and manufacturing (ISIC 1511-3720).[5] The data set comprises the assignment variable electricity use that determines if firms belong to treatment or control group, and five outcome variables. The outcome variables of our analysis are turnover, exports, investment, employment as measured by number of employees, and value added. Turnover, exports, investment, and value added are denoted in 1,000 euros. In addition, we show electricity intensity as descriptive statistic that is computed by dividing the amount of electricity use by turnover. The resulting index is denoted in KWh per euro.[6]

In Table 4.2, we present descriptive statistics for the original sample for the years 1995, 2000, and 2005.[7] Our data set includes close to 40,000 observations per year. As explained in Section 4.4.1, AFiD is a modular data set based on several different mandatory censuses and surveys. Hence, the sample size varies depending on the variable under investigation and the associated census or survey.[8] We have information on turnover,

[5] Regarding the classification by economic activity, we refer to the International Standard Industrial Classification of all economic activities (ISIC) Rev. 3.1, as adopted by the Statistical Commission of the United Nations.

[6] Electricity intensity may also be of interest as an outcome variable. Given its construction as electricity use over turnover and with electricity use being the assignment variable, it does, however, not provide any additional information to simply analyzing turnover. Figure 4.9 and Figure 4.10 in Appendix 4.A show the electricity intensity as function of electricity use for given years in order to shed some light on the previously described relationship.

[7] For all considered variables, outliers have been removed outside the 1^{st} and 99^{th} percentile.

[8] The characteristics turnover, exports, and employment are gathered monthly by the same census, the Monthly Report on Plant Operation. Investment and electricity use stem from different censuses, namely the Census on Investment, the Monthly Report on Plant Operation, and the Census on Energy Use. The Census on Investment is conducted yearly. While information on energy use was collected by the Monthly Report on a monthly basis from 1995 - 2002, an independent census on energy use was established in 2003.

Table 4.2: Descriptive statistics

	Mean	St. dev.	P10	P 50	P90	N
		1995				
Electricity use (in MWh)	1,346.66	3,474.06	37.40	284.92	3,170.90	38,470
Turnover (in EUR 1,000)	13,155.15	23,575.62	1,423.09	5,134.15	31,759.60	38,579
Exports (in EUR 1,000)	2,622.11	7,802.80	0	93.01	6,559.38	38,579
Investment (in EUR 1,000)	594.14	1,378.32	0	136.67	1,490.43	32,975
Employment	104.56	154.27	22.50	51.00	235.67	38,579
Electricity intensity (in KWh per EUR)	0.1003	0.1247	0.0110	0.0577	0.2414	38,470
Value added (in EUR 1,000)	-	-	-	-	-	-
		2000				
Electricity use (in MWh)	1,509.58	3,968.69	41.47	304.95	3,541.74	38,784
Turnover (in EUR 1,000)	14,855.25	27,579.86	1,520.13	5,462.99	36,230.26	38,873
Exports (in EUR 1,000)	3,726.30	11,062.76	0	129.68	9,378.87	38,873
Investment (in EUR 1,000)	603.73	1,423.36	0	135.71	1,518.55	36,493
Employment	99.81	141.20	22.75	49.5	228	38,873
Electricity intensity (in KWh per EUR)	0.1020	0.1262	0.0108	0.0599	0.2397	38,784
Value added (in EUR 1,000)	8,945.63	13,821.24	1,036.60	3,778.24	22,868.13	15,152
		2005				
Electricity use (in MWh)	1,888.30	4,938.04	60.51	400.43	4,437.14	36,158
Turnover (in EUR 1,000)	16,183.06	30,413.63	1,483.17	5,740.41	39,668.39	37,329
Exports (in EUR 1,000)	4,950.96	13,909.35	0	302.92	12,822.16	37,329
Investment (in EUR 1,000)	477.57	1,192.87	0	90.62	1,192.46	35,111
Employment	97.78	137.62	22.75	49.50	217.67	37,329
Electricity intensity (in KWh per EUR)	0.1201	0.1431	0.0144	0.0732	0.2773	35,897
Value added (in EUR 1,000)	9,502.641	14,542.27	1,039.019	4,089.146	24,673.86	13,997

Notes: Turnover, investment, and exports are denoted in EUR 1,000. Electricity use relates to the taxable electricity use in MWh (not including self-generated electricity). Electricity intensity is denoted by electricity use divided by turnover. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey, own calculations.

exports, investment, employment, and electricity use for all firms of the manufacturing sector with more than 20 employees summing up to about 40,000 observations on an annual basis from 1995 - 2005. For value added, we have only information from a random sample of about 15,000 firms on an annual basis from 1999 - 2005.

A comparison of the 10^{th} and 90^{th} percentile of the outcome variables and electricity use (Table 4.2) shows that the firms highly differ in their characteristics leading to high dispersion in the corresponding distributions. The percentiles as well as a comparison of mean and median show that the distributions of firms over the considered variables are positively skewed. This reflects the high fraction of small and medium sized firms and their importance for the German economy. About 90 percent of the firms in the census operate only a single plant.

Many firms operate around the thresholds for the reduced electricity rate, i.e., 50 MWh from 1999 to 25 MWh from 2003 onwards. Figure 4.2 shows histograms of the distribution of firms in the manufacturing sector ordered across their electricity use for the years 1995, 2000, and 2005. Each bin shows the absolute frequency of firms within the considered range. In the first row, we show histograms for the full support of electricity use, while the second row shows histograms with a trimmed support of 0 - 1,000 GWh. A bin corresponds to a 200 MWh range in the first row and a 20 MWh range in the second row. The histograms in the first row show very few firms with electricity use above 2,000 MWh while many more firms consume less than 2,000 MWh. The lowest bin in terms of electricity use, which corresponds to an electricity use of 0 to 200 MWh, contains close to 39.9 percent of all firms included in the data set in 2000. The histograms in the second row illustrate that there are more firms in the bins close to thresholds for a reduced electricity tax rate, i.e., around 50 MWh to 25 MWh, than in bins with higher electricity use. The high number of firms that consume less than 100 MWh enables us to perform the regression discontinuity analysis also for subpopulations, i.e. subsectors such as the manufacture of basic metal and metal products.

The corresponding Census on Energy Use collects information on energy use on a yearly basis from 2003 - 2005. Information about value added is collected by the annual Cost Structure Survey on a yearly basis from 1999 - 2005.

Figure 4.2: Histograms of electricity use in 1995, 2000, and 2005

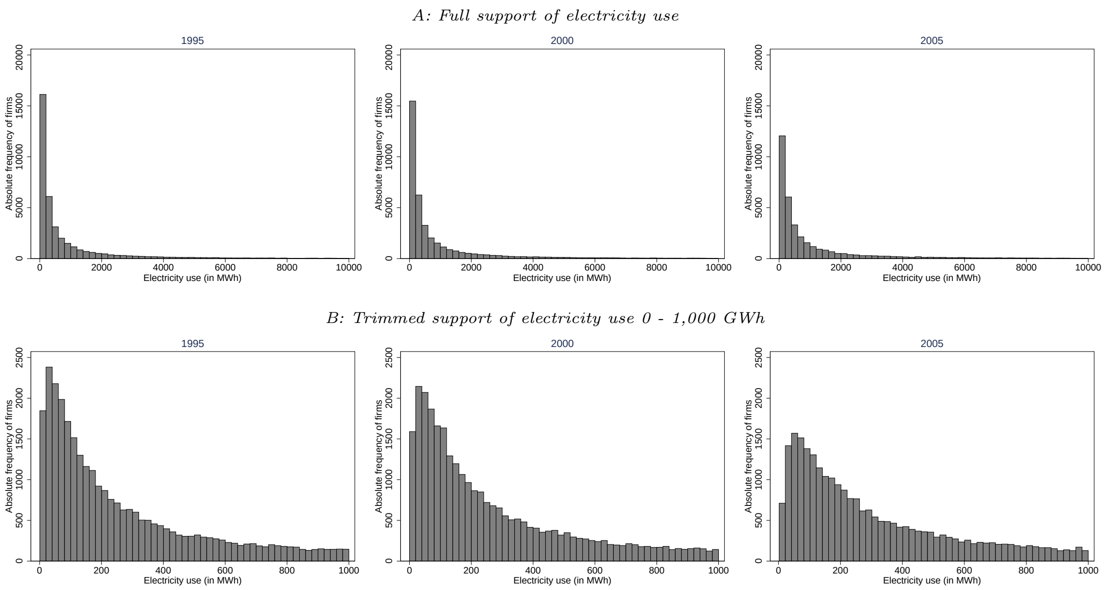

Notes: Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

4.5 Empirical evidence

4.5.1 Testing for identifying assumptions

In this section, we investigate the validity of our identification strategy. Applying the guideline set out by Lee and Lemieux (2010), we aim to confirm the assumptions that underlie the regression discontinuity design.

First, we examine the assumption that firms are unable to precisely control the assignment variable, i.e., electricity use. If this assumption holds, assignment to the treatment group is as good as locally random. According to Lee and Lemieux (2010) the incentive for sorting around the threshold is unproblematic, as long as the assignment variable contains a stochastic error component. In this case, optimizing firms do not have *precise* control over the assignment variable resulting into local random assignment to the treatment.

The assumption of imprecise control of the assignment variable cannot be directly tested. Nevertheless, by examining first the aggregate empirical distribution of the assignment variable and then applying a more formal test on the continuity of the distribution developed by McCrary (2008), we are able to shed light on the validity of this assumption.

In Figure 4.3, we present histograms that illustrate the distribution of the assignment variable electricity use for the pre-treatment year 1995 and the treatment years 1999 - 2005. The support of each distribution is trimmed to a range of 100 MWh. The graphs show the absolute frequencies of firms computed over non-overlapping bins with a bandwidth of 1 MWh. Following Lee and Lemieux (2010), we choose binwidths as small as possible, that still allow us to see the shape of the distribution. The vertical black line in each graph denotes the threshold at which the marginal tax rate changes in that year (the graph illustrating the pre-treatment year 1995 shows the threshold of 1999).

The bin-to-bin jumps in the frequencies enable us to identify exceptional jumps at the threshold c that indicate a discontinuity in the density. If firms could precisely manipulate their electricity use and thereby select themselves into the treatment group, we would expect a significant upward jump in the bins located directly right of the threshold.

The histograms do not provide any evidence that firms manipulated their electricity use. Figure 4.3 shows several upward jumps that are located far from the thresholds. However, directly right of the thresholds there are no unusual jumps that would indicate manipulation of electricity use.

Figure 4.4 shows a visualization of the discontinuity test developed by McCrary (2008) for the pre-treatment year 1995 and the treatment years 1999 - 2005. Each graph exhibits an estimate of the density function of the assignment variable electricity use and the corresponding 95 percent confidence interval. The density function is estimated using the local linear density estimation technique proposed by McCrary (2008). The dots represent local densities for bins with a width between 0.50 and 0.75 MWh. The binwidths are calculated following the procedure in McCrary (2008).

Examining the point cloud, which gives a good visual impression of the empirical density function of the assignment variable, we do not see clear evidence for a discontinuity at the threshold in the pre-treatment year 1995 and in the treatment years 1999 - 2004. An inspection of the plotted density function and the corresponding confidence intervals lead to the same result. Only for the year 2005, the test shows that the density is significantly higher close to the right of the threshold suggesting a discontinuity at the threshold. Yet, looking at the absolute frequencies for the same year in Figure 4.3 also reveals excess mass close to the left of the threshold. In particular, the number of firms increases sharply at 24 MWh electricity use. In comparison to the jumps and irregularities in the absolute frequencies further away from the threshold, there is a slight unsystematic jump in the

Figure 4.3: Distribution of electricity use near the threshold in 1995 and 1999 - 2005

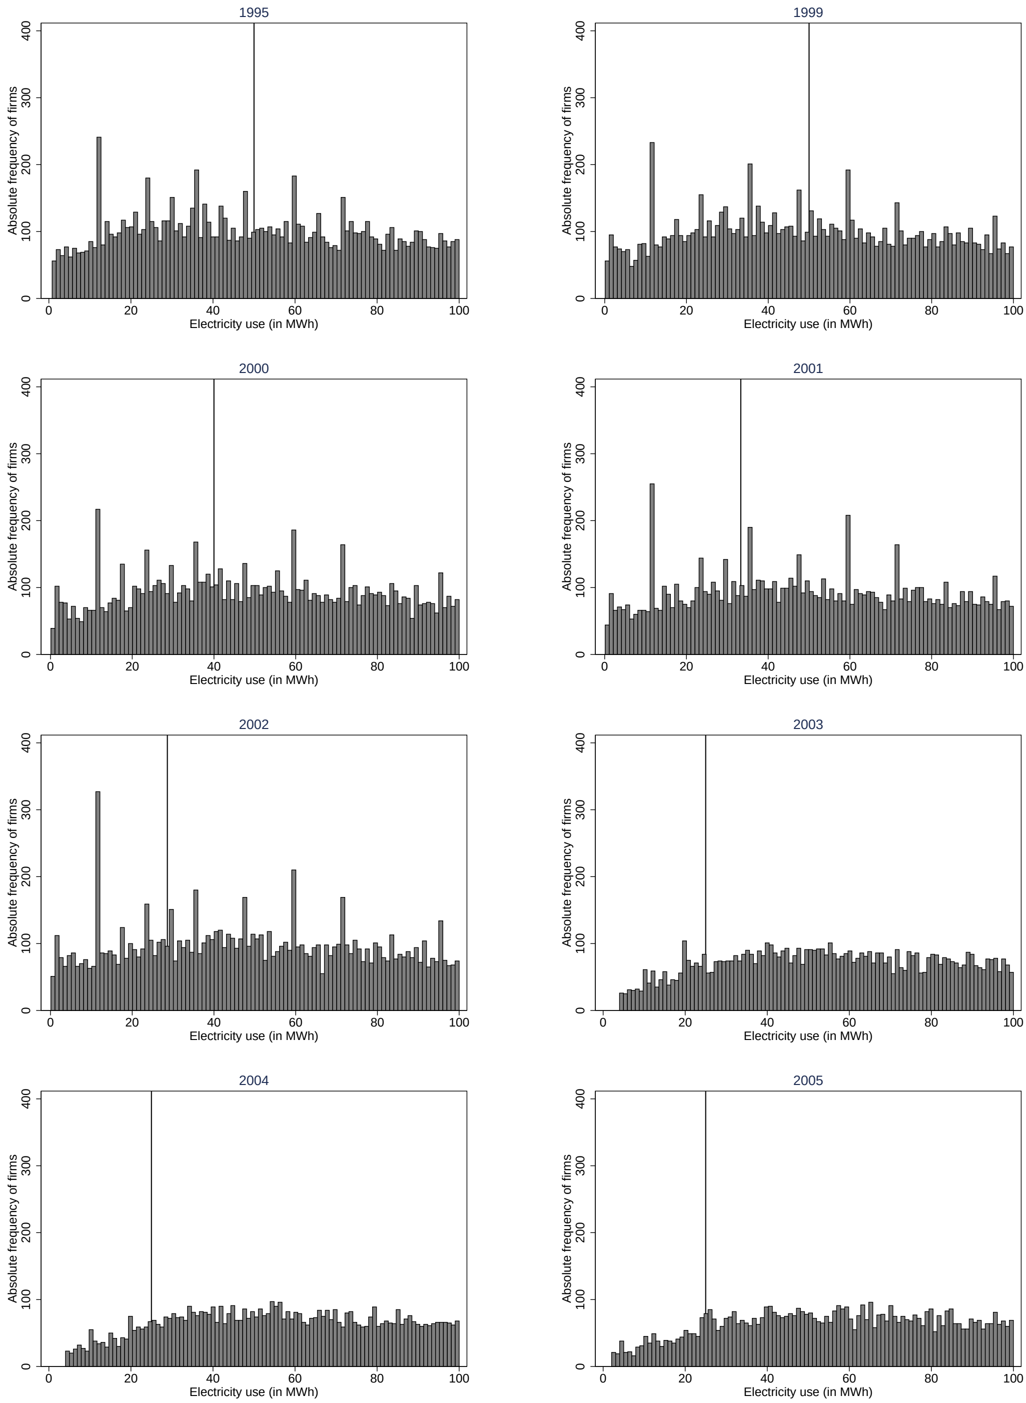

Notes: Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

Figure 4.4: Visualization of the McCrary discontinuity test

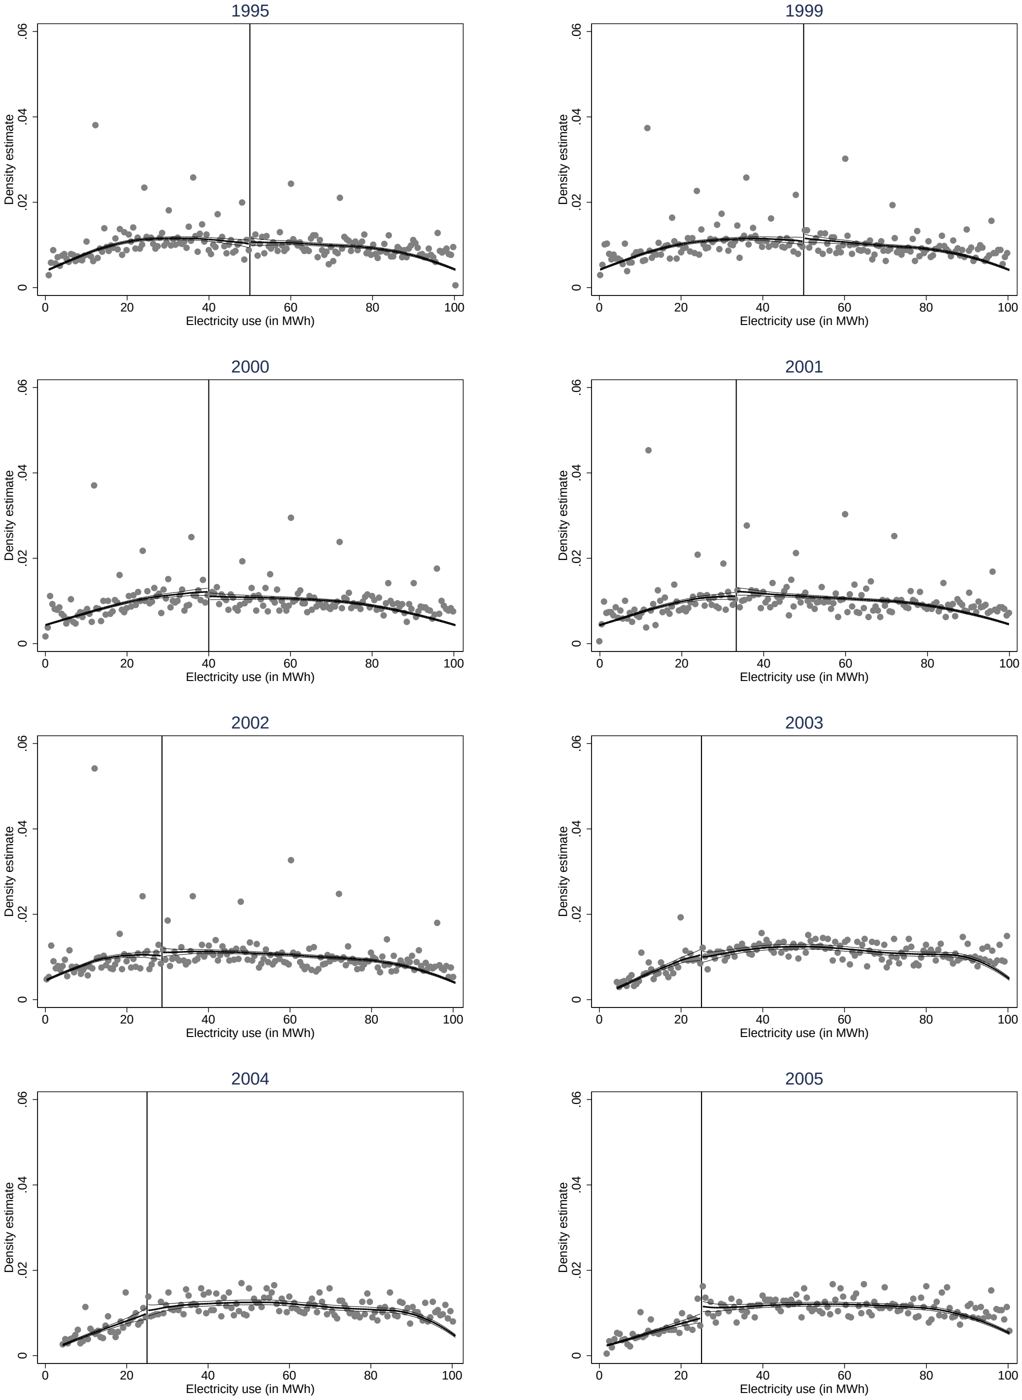

Notes: Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

frequencies between 24 and 27 MWh electricity use. The rejection of the null hypothesis of continuity in the framework of the test developed by McCrary (2008) may therefore be due to an unsystematic jump in the density rather than a systematic break in the density function. Also, the graphs in Figure 4.4 show jumps in the local densities for all years, even at locations far away from the thresholds.

An alternative approach for investigating a potential sorting into the treatment group would be to examine, if continuous baseline covariates show discontinuities at the threshold. However, for firm data, this approach is barely feasible, since one would need firm characteristics that are (i) continuous and (ii) unaffected by the treatment. A change in the relative input prices - e.g. through a tax - potentially leads to a change in input use as well as output production. All continuous variables in our data set hence might be affected by the electricity tax.

From 2003 onwards, the histograms as well as the density estimates show fewer firms in comparison to the years before. This phenomenon emerges mostly due to two methodological changes. First, due to a switch from the monthly to the yearly census, some firms were not surveyed in the years 2003 and 2004. Second, the classification by economic activity changed in 2003. Firms may have ascertained, whether they were correctly classified. Consequently, some firms that actually were not in the manufacturing sector might have been reclassified and disappeared from the data set.

The second assumption, we investigate is the assumption of local continuity. In particular, we assume that the outcome variables evolve as continuous functions of electricity use around the threshold when the intervention is absent. Since we do not observe the counterfactual - i.e. firms that lie above the threshold and are not treated - we analyze the relationship of outcome and assignment variable before the intervention started. Figure 4.5 contains four scatter plots showing the outcome variables turnover, exports, investment, and employment as second order global polynomial functions of electricity use for the pre-treatment year 1995. The dots denote non-overlapping binned local means of the corresponding outcome variable. The local means are computed for 1 MWh bandwidths in the area of 25 - 75 MWh, the $c \pm 25$ MWh neighborhood of the 50 MWh threshold that applies for the first year of the treatment 1999. Neither the point cloud of binned local means, nor the second order polynomial give rise to concern that a discontinuity and thus a violation of the local continuity restriction is present.

Figure 4.5: Outcomes in the pre-treatment year 1995

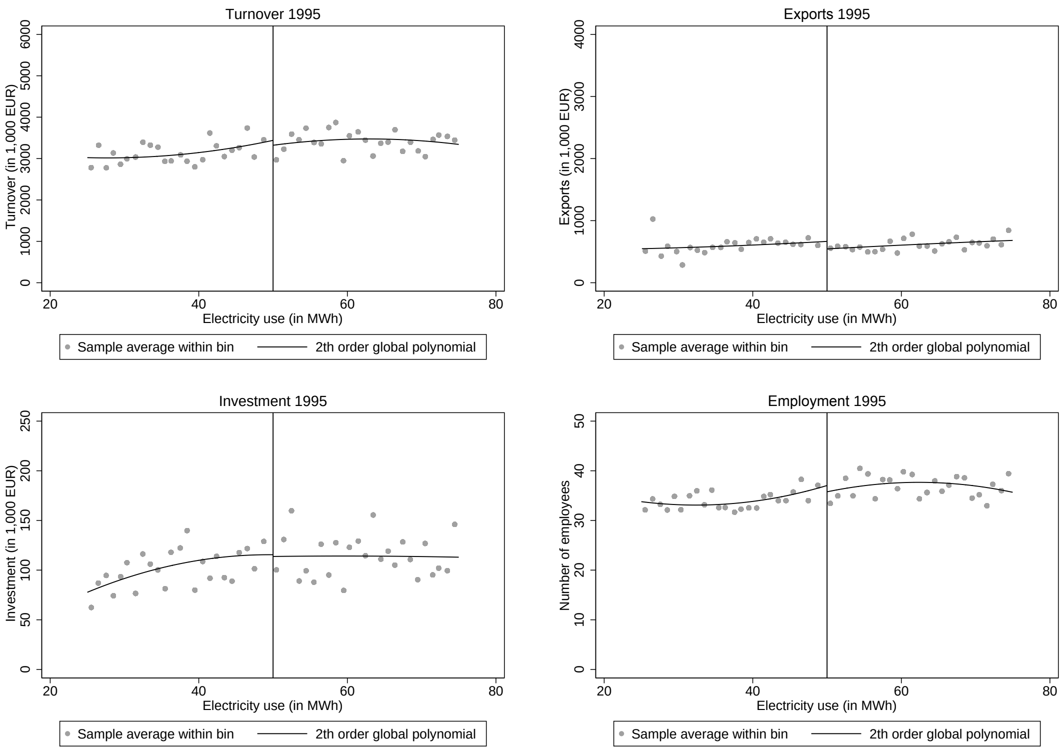

Notes: Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

4.5.2 Graphical analysis

We start our analysis by showing graphical evidence on the relationship between the outcome variables and the assignment variable electricity use. We compute local conditional sample averages for 1 MWh non-overlapping bins of electricity use and also show estimates of second order global polynomial regression functions for either side of the threshold separately. The graphs in the first column in Figure 4.6 show the results for four outcome variables: turnover, exports, investment, and employment in 2000. The graphs in the second column show the results for the same variables in 2005. The vertical black lines at 40 MWh and 25 MWh denote the thresholds for tax reductions. The plots are trimmed to the electricity use $c \pm 25$ MWh around the threshold.

Our aim is to discover discontinuities (or in other words shifts) in the local conditional sample averages. A shift at the threshold would indicate an effect of the tax reduction on the outcome variables. Shifts in regions away from the threshold would highlight the presence of other discontinuities and would question the applicability of the regression discontinuity design in this context. Note that the cloud of local conditional sample averages indicates the level of dispersion of the data.

The graphs depicted in Figure 4.6 do not show evidence of an obvious discontinuity at the threshold. A positive effect of the reduced tax rate on one of the outcome variables would be indicated by an upward shift to the right of the thresholds of both the binned averages and the regression lines. A negative effect on one of the outcome variables would be indicated by a downward shift to the right of the threshold of both the binned averages and the regression lines. Regarding the global polynomial functions, one should bear in mind that the estimates are less precise close to the thresholds than further away from them. A point estimated further away from the threshold can draw on additional information toward its right and left for estimation, while a point close to the threshold can only draw on additional information on one side. The small discontinuities in regression lines are thus likely due to less precise estimation at the thresholds than further away.[9]

[9] For the estimation of the local average treatment effect, in the following section, we rely on the nonparametric approach based on weighted local linear regressions on both sides of the threshold proposed by Hahn, Todd, and van der Klaauw (2001) and Porter (2003) in order to mitigate this problem. The estimator shows good small sample properties and is suitable for inference at the boundary of the support of the regression function.

Figure 4.6: Effects caused by the discontinuity (I/II)

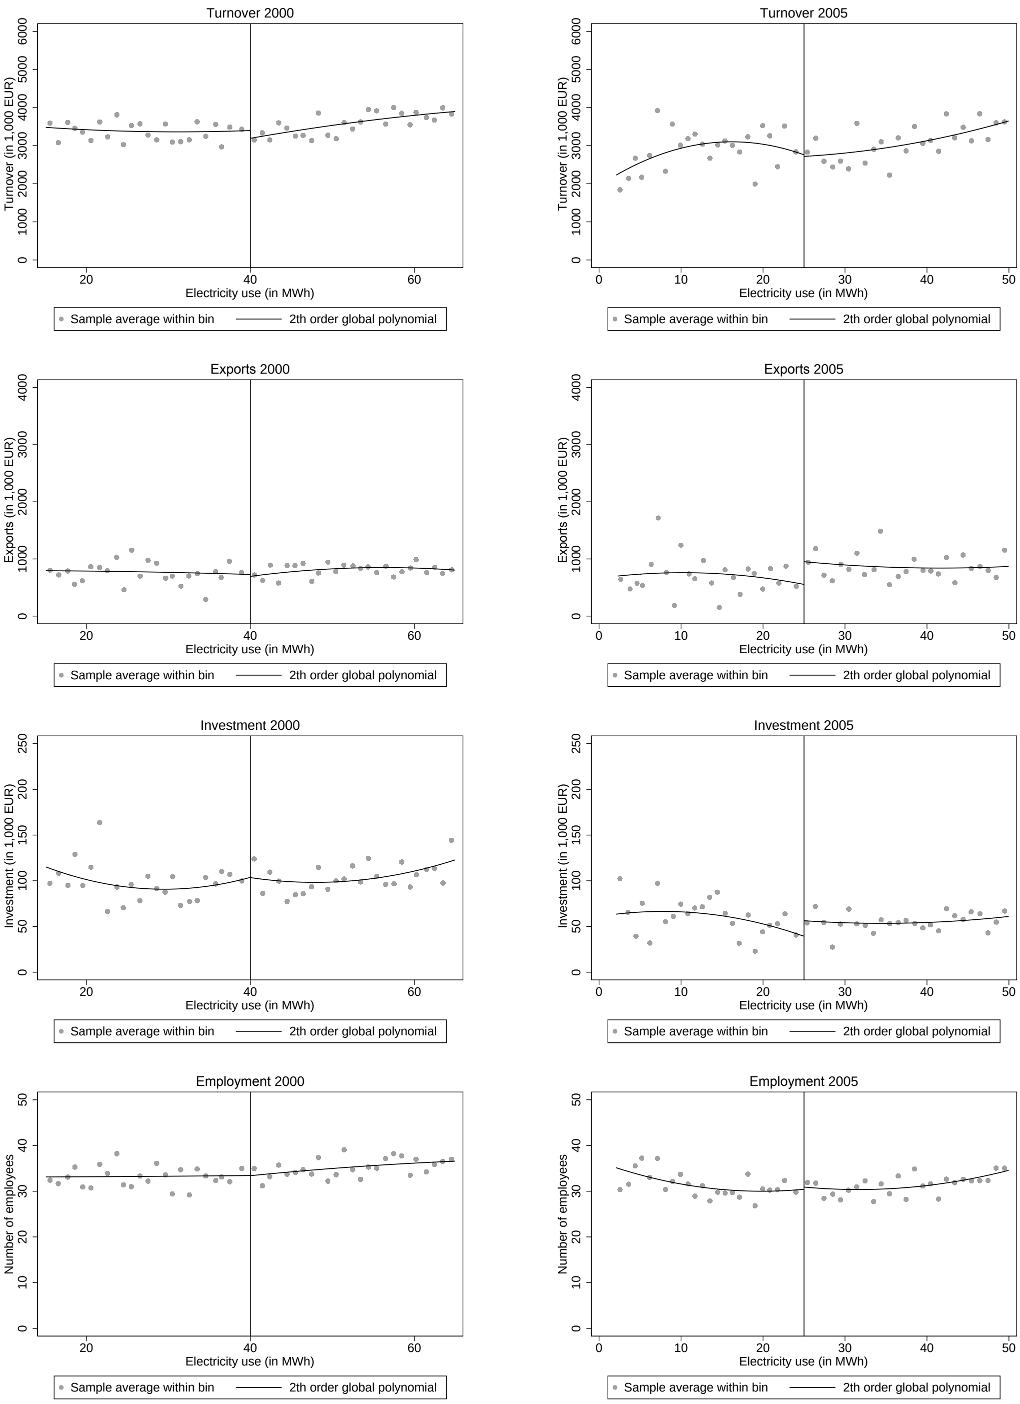

Notes: Assignment variable: electricity use. Outcome variables: turnover, exports, investment, and employment. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

A systematic shift of the regression lines or the cloud of binned local means indicating a discontinuity at the threshold is not observed.

For both turnover and investment, substantial heterogeneity is observed between the local sample averages reflecting the high degree of variance discussed in Section 4.4.3. However, no discontinuity is found at the threshold. Also for exports, the local sample averages do not indicate a discontinuity at the threshold. However, the global polynomial function indicates a slight upward shift to the right of the threshold. This is seemingly driven by the four bins to the left of the threshold for the reduced tax rate at 25 MWh and the five bins to the right of the threshold. Bins further away from the threshold do not show a consistent difference in average exports. No indication for a discontinuity at the threshold is found for employment, neither by the local sample averages nor by the global polynomial functions.[10]

Figure 4.7 shows the impact of the reduced electricity tax on value added. Information on value added is only available from a mandatory survey of a subset of firms. Therefore there are less observations than for the outcome variables above that originate from the census of firms. The dispersion of value added is lower than that of turnover or exports as also shown in the descriptive statistics in Section 4.4.3. This translates into a fairly smooth relationship between value added and electricity use and may help to detect a potential discontinuity at the threshold. However, neither the binned conditional sample averages nor the global polynomial regression functions indicate an effect of the reduced electricity tax on value added.

In addition, the plots do not provide evidence for discontinuities away from the threshold. Hence there is no indication of other sources that may cause discontinuities in the relationships between outcome variables and assignment variable.

[10]The observed pattern for the years 2000 and 2005 also holds for other years in which the reduced tax rate applied. Results are available upon request.

Figure 4.7: Effects caused by the discontinuity (II/II)

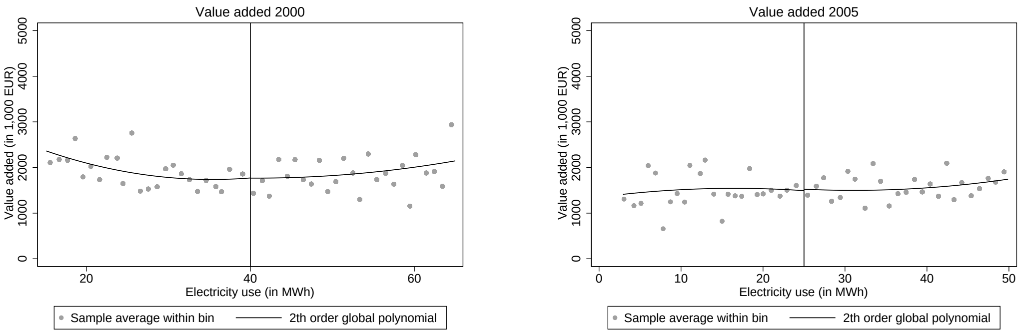

Notes: Assignment variable: electricity use. Outcome variable: value added. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey, own calculations.

4.5.3 Local average treatment effects

In this section we present the estimated local average treatment effects of the tax reduction scheme on the outcome variables turnover, exports, investment, employment, and value added. Being precise, we estimate the effect of the difference between the full and the reduced marginal tax rate - i.e. the reduction of the marginal tax rate. The firms that consume more electricity than the threshold c benefit from a lower marginal tax rate and form the treatment group. The firms that consume less electricity than the threshold c face the full marginal tax rate and thus denote the control group. A year by year evaluation leads to seven experiments and 35 treatment effects of interest in the years 1999-2005.

The estimators of the local average treatment effects presented in the following are computed as described in Section 4.3.3. Recall, that the estimators of the local average treatment effects are computed as the difference of the constants of two weighted linear regressions for narrow bandwidths left and right of the threshold. Here, the weights for linear regression are computed based on a triangular kernel function.[11] The bandwidths are computed based on the data-driven bandwidth selection procedure developed by Imbens and Kalyanaraman (2012).

In Table 4.3, we show the estimated effects of the tax reduction for each year in the

[11]The results do not systematically change when using alternative kernel functions. Table 4.9 and 4.10 in Appendix 4.B report the results of the local average treatment effect estimation considering uniform and Epanechnikov kernel functions.)

Table 4.3: Local average treatment effects

	Tax reduction scheme			Effect of reduced marginal tax rate				
Year	Threshold (MWh)	Full tax rate (EUR/MWh)	Tax reduction (EUR/MWh)	Turnover	Exports	Investment	Employment	Value added
1999	50	10	8	95.40	2.01	-10.50	-0.39	-83.75
				(169.85)	(108.37)	(11.24)	(0.99)	(199.12)
2000	40	12.5	10	-166.78	-36.53	-1.73	-0.12	-18.67
				(180.53)	(108.98)	(11.54)	(1.17)	(200.28)
2001	33	15	12	440.78*	-180.18	9.36	-0.62	183.14
				(216.96)	(121.50)	(9.80)	(0.96)	(208.51)
2002	28.6	17.9	14.6	-379.65	-47.27	-20.65*	0.16	-492.54
				(238.68)	(108.33)	(10.29)	(1.13)	(299.71)
2003	25	20.5	8.2	-136.42	-232.44	-4.18	-0.49	-177.09
				(221.77)	(156.74)	(8.43)	(1.33)	(181.25)
2004	25	20.5	8.2	254.35	-48.75	-4.41	0.72	83.51
				(216.70)	(157.89)	(9.00)	(1.04)	(203.20)
2005	25	20.5	8.2	-106.73	335.86*	14.48	0.59	-35.59
				(268.37)	(164.23)	(7.88)	(1.32)	(213.67)

Notes: * indicates significance at the 5 percent level. Standard errors are shown in parentheses. Analysis covers firms in the ± 25 MWh region around the threshold. The order of the polynomial function is set to 1. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey, own calculations.

period 1999 - 2005 along with the characteristics of the prevailing tax scheme. In each experiment we consider observations in the neighborhood $c \pm 25$ MWh around the threshold. Outliers of the outcome variables are removed outside the 1^{st} and 99^{th} percentile. The columns on the left side of Table 4.3 summarize the information on the electricity tax. They show for each year the full tax rate as well as the thresholds from which the reduced marginal tax rate applies and the difference between the full marginal tax rate. The columns on the right side of the table show the point estimates of the regression discontinuity analysis and the corresponding standard errors.

The thirty-two statistically insignificant effects in Table 4.3 clearly outweigh the three statistically significant effects. These statistically significant effects indicate a positive impact of the tax reduction on turnover in 2001 and exports in 2005 as well as a negative effect on investment in 2002.

Table 4.4 shows the bandwidth choice for each experiment as well as the number of observations that lie within the bandwidths right and left of the threshold. The selected bandwidths lie in a range between 15 and 25 MWh. The selected bandwidths for exports are typically smaller and thereby have fewer observations than those for turnover, investment, and employment.

The results from the regression discontinuity analysis indicate hardly any evidence for a consistent effect of the reduced marginal electricity tax on turnover, exports, investment, employment, or gross value added. First, there is only a low number of statistically significant treatment effects (only three out of thirty-five) that might result from statistical error. Second, there is no consistent pattern of negative or positive signs for the local treatment effects. Neither do the three statistically significant effects have the same sign nor do the five dependent variables show a particular pattern or trend.[12]

[12] To investigate robustness with regard to distributional assumptions and outliers, we estimated a log specification of our model. The qualitative findings do not change. Detailed results are available upon request.

Table 4.4: Imbens and Kalyanaraman (2012) bandwidths and number of observations

Outcome variable	Bandwidth	Number of observations		
		c ± 25 MWh	Control group	Treatment group
		A: 1999		
Turnover	24.15	5,289	2,671	2,442
Exports	16.11	2,330	755	793
Investment	22.89	3,873	1,848	1,739
Employment	23.42	5,289	2,615	2,377
Value added	21.30	1,452	661	600
		B: 2000		
Turnover	22.38	5,017	2,306	2,263
Exports	18.47	2,137	772	815
Investment	19.07	3,691	1,487	1,397
Employment	19.34	5,014	2,023	1,877
Value added	20.17	1,301	536	546
		C: 2001		
Turnover	16.61	4,862	1,557	1,769
Exports	18.03	2,041	647	842
Investment	17.48	3,338	1,095	1,302
Employment	25.00	4,859	2,339	2,520
Value added	20.35	1,119	413	495
		D: 2002		
Turnover	14.01	5,072	1,323	1,511
Exports	18.07	2,114	758	819
Investment	20.85	3,360	1,316	1,572
Employment	20.32	5,063	2,047	2,216
Value added	22.37	985	377	510
		E: 2003		
Turnover	16.28	3,052	891	1,294
Exports	12.74	1,290	278	407
Investment	18.35	2,175	650	1,076
Employment	18.97	3,052	964	1,537
Value added	16.36	851	249	362
		F: 2004		
Turnover	14.08	2,779	657	1,079
Exports	14.44	1,138	236	466
Investment	17.44	1,979	553	960
Employment	18.44	2,778	798	1,414
Value added	15.82	704	172	319
		G: 2005		
Turnover	12.22	2,654	535	843
Exports	17.12	1,068	266	479
Investment	17.36	1,856	495	870
Employment	12.78	2,654	559	886
Value added	23.17	621	177	408

Notes: Turnover, investment, and exports are denoted in EUR 1,000. The number of observations refers to the ± 25 MWh region around the threshold c. The bandwidth is selected based on Imbens and Kalyanaraman (2012). Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey, own calculations.

4.5.4 Sensitivity toward bandwidth choice

In this section, we investigate the sensitivity of our findings toward different bandwidths. The results in the previous section do not show any systematic significant effects of the reduced tax rates on economic outcomes. The question in the following paragraphs is whether these results are robust for various choices of bandwidth.

Bandwidth choice is a choice between precision and bias. Larger bandwidths offer more precise estimates as they can rely on a larger number of observations. At the same time, larger bandwidths may generate bias, in particular, when using a linear estimator for data that is inherently nonlinear. The optimal bandwidth that minimizes the mean squared error decreases with the number of observations. In the previous section, we selected bandwidths according to a fully data driven and asymptotically optimal bandwidth choice algorithm developed by Imbens and Kalyanamaran (2012).

Given the above mentioned tradeoffs between precision and bias, we present results across different integer bandwidth choices ranging from 5 to 25 MWh in Figure 4.8 for the years 2000 and 2005. The solid black line in each graph denotes point estimates and the dashed lines are corresponding 95 percent confidence intervals. The standard errors decrease with increasing bandwidths as expected. In most cases, also the estimates become smaller in absolute terms and approach zero with increasing bandwidths, without becoming statistically significant. This confirms our previous findings that do not indicate any effects of the reduced tax rates on economic outcomes. Smaller bandwidths tend to have larger point estimates. Given the higher imprecision of the estimates, no point estimate is significant for bandwidths below 16 MWh, adding to the evidence that there is no significant effect.

In addition, we note that the observed patterns for 2000 and 2005 hold for the other years too. Table 4.11 in Appendix 4.C reports the results of the local average treatment estimation for the bandwidths 5, 10, 15, 20, and 25 in 1999-2005. The significant positive local average treatment effect on turnover in 2001 does not seem to depend on bandwidth choice. Yet, the significant negative estimate for investment in 2002 is not robust to bandwidth choice. It is only significant for a bandwidth between 10 and 20 MWh. Figure 4.11 in Appendix 4.C shows the point estimates and 95 percent confidence intervals across different integer bandwidth choices for turnover in 2001 and investment in 2002.

Figure 4.8: The effects of bandwidth choice on point estimates and confidence intervals

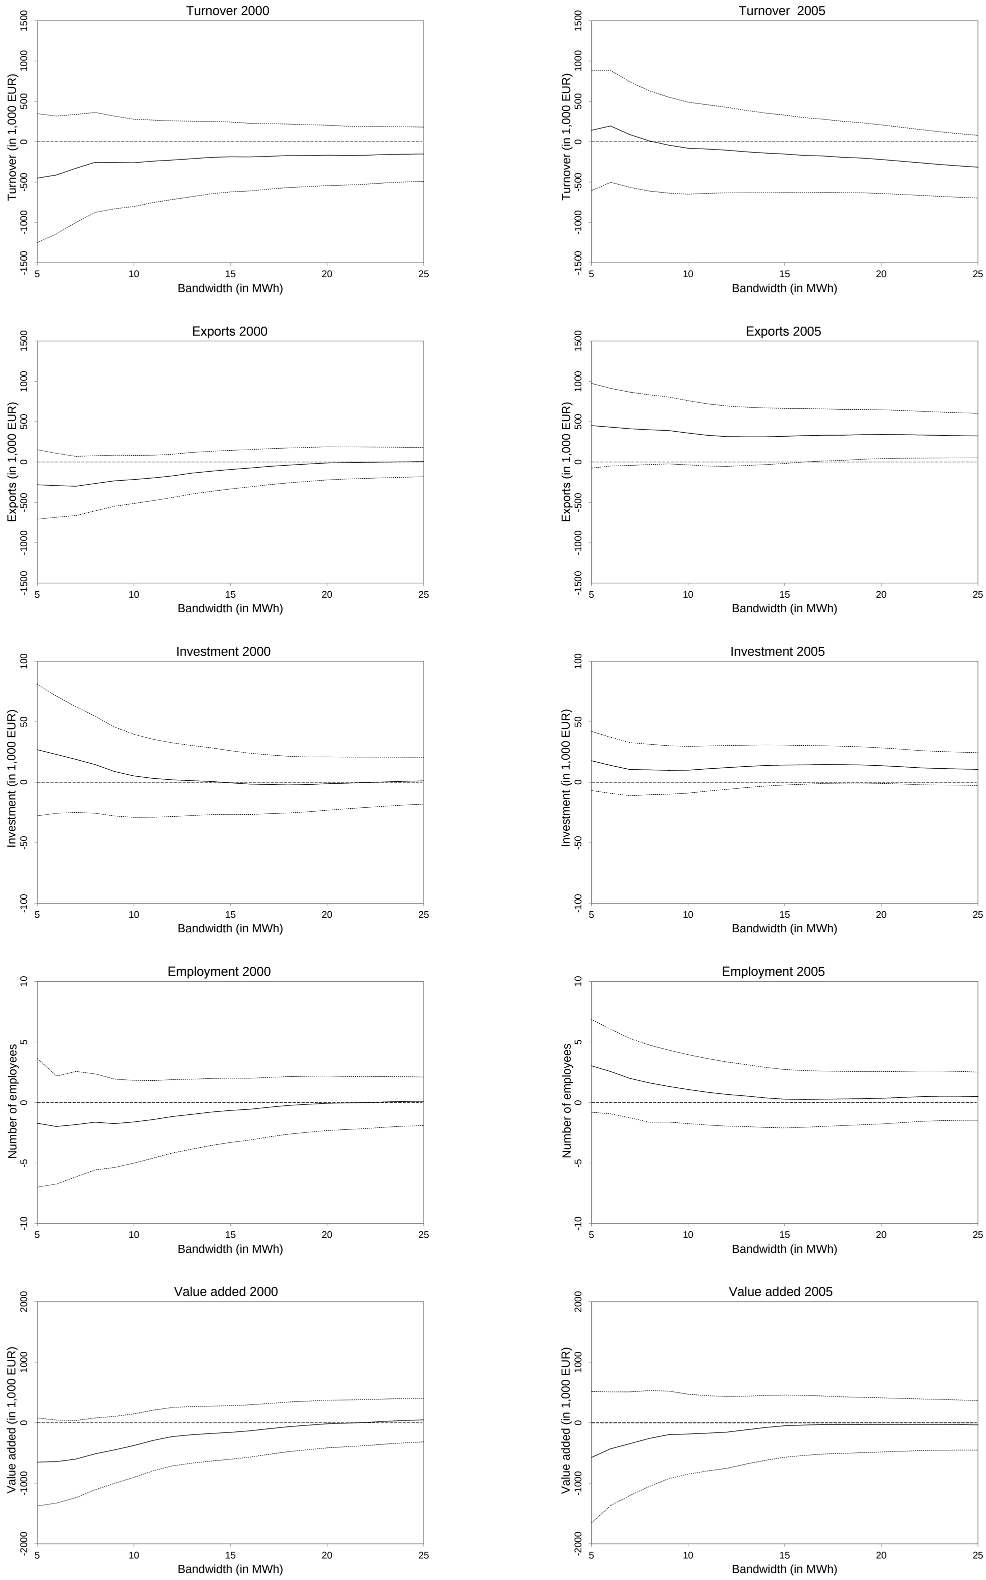

Notes: The solid black line in each graph denotes point estimates and the dashed lines are corresponding 95 percent confidence intervals. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey, own calculations.

4.5.5 Sensitivity toward polynomial choice

In addition to selecting the bandwidth, the choice of the polynomial order may also affect results. Choosing a local linear estimator for data that is inherently non-linear may bias results, in particular when the bandwidth is large. While Figure 4.6 in Section 4.5.2 might suggest that higher order global polynomial estimators fit best for some outcome variables in some years, it also does not point toward strong local non-linearities in the data. This visual inspection may therefore suggest that the previously chosen local linear regressions should not suffer from substantial bias.

An additional robustness check with a higher order polynomial does not change the previous findings, further confirming that the local linear regressions are not substantially biased. Table 4.5 shows the results for the effect of the reduced tax rate on economic outcomes applying local quadratic polynomial regressions. The bandwidths are optimally selected using the algorithm developed by applying Imbens and Kalyanamaran (2012) as previously. While many point estimates increase somewhat confidence intervals also increase substantially. This results in only three out of 35 estimates becoming statistically significant. No pattern regarding the signs of the effects is observed, confirming that there are no consistent effects of the reduced tax rates on economic outcomes.

Given the fairly linear underlying data close to the threshold, results would unlikely change with higher polynomial orders. The underlying data nevertheless reveals a fair amount of heterogeneity as both shown in Figure 4.6 and the descriptive statistics. The following section therefore investigates how this heterogeneity may impact our results.

Table 4.5: The effects of an alternative polynomial order

Outcome variable	Estimator	Bandwidth	Number of observations		
			c ± 25 MWh	Control group	Treatment group
A: 1999					
Turnover	163.50 (238.29)	23.43	5,289	2,606	2,387
Exports	-77.66 (144.47)	19.93	2,330	901	965
Investment	-9.45 (15.91)	23.41	3,873	1,883	1,770
Employment	-0.62 (1.41)	21.19	5,289	2,375	2,124
Value added	-157.62 (291.27)	18.83	1,452	589	540
B: 2000					
Turnover	-241.09 (282.45)	20.39	5,017	2,118	2,069
Exports	-196.71 (153.00)	21.03	2,137	854	972
Investment	-0.93 (17.01)	21.72	3,691	1,615	1,636
Employment	-1.72 (1.86)	18.75	5,014	1,976	1,834
Value added	-331.80 (272.04)	21.07	1,301	550	568
C: 2001					
Turnover	580.77* (286.41)	21.22	4,862	1,915	2,197
Exports	-182.94 (177.07)	17.71	2,041	639	825
Investment	7.76 (12.49)	21.91	3,338	1,437	1,607
Employment	0.30 (1.49)	22.27	4,859	2,164	2,294
Value added	-368.42 (272.01)	21.87	1,119	466	532
D: 2002					
Turnover	-430.40 (266.11)	24.35	5,072	2,335	2,628
Exports	-65.91 (135.57)	29.34	2,114	910	1,204
Investment	-27.84* (13.49)	28.42	3,360	1,514	1,846
Employment	-0.63 (1.58)	23.32	5,063	2,264	2,535
Value added	-911.67* (448.78)	21.69	985	367	495
E: 2003					
Turnover	41.90 (323.79)	16.66	3,052	899	1,329
Exports	-357.13 (215.91)	15.66	1,290	333	529
Investment	2.08 (11.82)	19.23	2,175	664	1,136
Employment	1.16 (2.12)	16.60	3,052	888	1,330
Value added	-134.90 (269.32)	14.58	851	232	318
F: 2004					
Turnover	420.63 (274.62)	17.46	2,779	776	1,345
Exports	-112.14 (231.50)	15.07	1,138	251	490
Investment	-1.15 (12.77)	20.36	1,979	600	1,120
Employment	2.81 (1.46)	18.53	2,778	802	1,420
Value added	107.21 (292.40)	21.79	704	202	441
G: 2005					
Turnover	-29.45 (319.81)	17.54	2,654	718	1,256
Exports	334.83 (224.97)	17.00	1,068	264	473
Investment	11.09 (10.77)	17.80	1,856	498	891
Employment	0.33 (1.47)	21.73	2,654	808	1,566
Value added	-167.17 (355.19)	20.39	621	170	354

Notes: * indicates significance at the 5 percent level. Standard errors are shown in parentheses. Turnover, investment, and exports are denoted in EUR 1,000. The number of observations refers to the ± 25 MWh region around the threshold c. The bandwidth is selected based on Imbens and Kalyanaraman (2012). Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey own calculations.

4.5.6 Treatment effects across industries

The aim of looking at treatment effects across industries is twofold. First, we shed light onto the robustness of our results with respect to heterogeneity across industries within the manufacturing sector. For this purpose, we analyze the effect of the electricity tax reduction on firms of different subpopulations. Second, we aim to examine the effect of the electricity tax reduction on an energy intensive industry. If the electricity tax reduction has no impact on firms of an industry that is particularly affected by higher electricity prices, this would add additional support to the findings in the previous sections.

Industries within manufacturing differ along many dimensions. These differences concern - among other things - the output they produce, the technologies they deploy, or the market and industry structures they face. As a consequence, the treatment effect of the electricity tax reduction may vary across industries or subsectors. If effects have different signs for different subpopulations, this might lead to an insignificant average treatment effect for the whole population. In addition, if effects are only significant for a small subpopulation that is very energy intensive, this might not show up in the average treatment effect for the whole population of firms.

The first subpopulation are firms that manufacture machinery, electronic devices, and vehicles.[13] This subpopulation is chosen, first, as it consist of more homogenous firms compared to firms that produce all other types of goods and second, it still comprises a sufficient number of firms to conduct a regression discontinuity analysis. On average, this group shows higher turnover and exports and is less energy intensive in comparison to the full population.[14]

Table 4.6 shows the estimated effects of the tax reduction on the outcome variables turnover, exports, investment, and employment for each year in the period 1999-2005. We cannot estimate the effects on value added given too few observations from the sampled Cost Structure Survey. We apply local linear regressions and choose bandwidths optimally

[13]According to ISIC Rev. 3.1: manufacture of machinery and equipment n.e.c. (29), manufacture of office, accounting, and computing machinery (30), manufacture of electrical machinery and apparatus n.e.c. (31), manufacture of radio, television, and communication equipment and apparatus (32), manufacture of medical, precision and optical instruments, watches and clocks (33), manufacture of motor vehicles, trailers, and semi-trailers (34), manufacture of other transport equipment (35)

[14]In Appendix 4.E, Table 4.13 and 4.14 show detailed descriptive statistics of the two sub populations under investigation.

selected by applying Imbens and Kalyanamaran (2012) as in Section 4.5.3. The results do not provide evidence for a significant and systematic effect of the electricity tax reduction on the outcome variables. Only one out of thirty-five treatment effects is statistically significant. As for the whole population, the results show a significant effect on exports in 2005.

The second subpopulation are firms that manufacture basic metals and fabricated metal products.[15] The manufacturing of metals is a very energy-intensive manufacturing sector. This group should therefore be more sensitive with regard to changes in electricity prices. On average, firms of this group use higher amounts of electricity per unit of output, produce less output in terms of turnover and export less than the average firm of the full population. Table 4.7 shows the treatment effect for manufacture of basic metals and fabricated metal products. The local treatment effect of the electricity tax reduction on turnover is significantly positive in 2005. All other effects are statistically insignificant. Even for this more homogenous and energy intensive sector, we do not find evidence for a significant and systematic effect of the electricity tax reduction.

The point estimates in both subpopulation analyses do not differ systematically from the point estimates of the analysis on the whole population. Hence, we do not observe any trend in the size of effects within the subpopulations, as may have been expected for the more energy-intensive manufacturing of metals. The standard errors of the subpopulation analyses are larger compared to whole population. This decrease in precision can be explained by the lower number of observations.

[15] According to ISIC Rev. 3.1: manufacture of basic metals (27) and manufacture of fabricated metal products, except machinery and equipment (28)

Table 4.6: Subsample analysis: manufacture of machinery, electronic devices, and vehicles

Outcome variable	Estimator	Bandwidth	Number of observations		
			c ± 25 MWh	Control group	Treatment group
			A: 1999		
Turnover	31.18 (348.01)	16.60	2,078	740	666
Exports	-38.12 (162.85)	24.40	1,139	538	584
Investment	-0.47 (18.36)	19.71	1,628	669	613
Employment	-0.84 (1.73)	16.31	2,078	727	645
			B: 2000		
Turnover	-203.59 (293.47)	23.29	1,986	925	956
Exports	-254.12 (191.07)	18.74	1,067	372	423
Investment	-15.14 (26.36)	13.51	1,570	447	446
Employment	-2.77 (2.04)	17.85	1,986	742	703
			C: 2001		
Turnover	-69.31 (354.97)	15.47	1,939	590	699
Exports	-292.90 (229.47)	14.99	997	264	364
Investment	-3.04 (15.82)	15.35	1,473	447	534
Employment	-2.89 (2.46)	14.02	1,939	532	616
			D: 2002		
Turnover	-496.64 (332.78)	18.16	2,095	749	847
Exports	-220.27 (224.59)	14.61	1,061	269	348
Investment	-25.61 (16.73)	16.29	1,522	425	571
Employment	1.14 (2.27)	20.07	2,094	794	960
			E: 2003		
Turnover	-267.88 (413.95)	13.38	1,342	328	443
Exports	-68.16 (264.04)	15.50	677	176	260
Investment	7.37 (13.05)	15.45	1,008	274	391
Employment	1.50 (3.14)	12.02	1,341	292	393
			F: 2004		
Turnover	255.51 (318.39)	15.64	1,273	319	554
Exports	32.12 (336.48)	14.92	611	130	242
Investment	-6.62 (11.40)	14.92	981	222	404
Employment	0.31 (1.79)	20.15	1,273	362	718
			G: 2005		
Turnover	684.84 (423.16)	9.13	1,253	193	312
Exports	682.02* (307.40)	16.47	566	129	237
Investment	8.81 (10.94)	20.00	947	269	504
Employment	1.62 (2.20)	15.36	1,253	300	518

Notes: * indicates significance at the 5 percent level. Standard errors are shown in parentheses. Turnover, investment, and exports are denoted in EUR 1,000. The number of observations refer to the ± 25 MWh region around the threshold c. The bandwidth is selected based on the procedure in Imbens and Kalyanaraman (2012). Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey, own calculations.

Table 4.7: Subsample analysis: manufacture of basic metal and fabricated metal poducts

Outcome variable	Estimator	Bandwidth	Number of observations		
			c ± 25 MWh	Control group	Treatment group
		A: 1999			
Turnover	-49.29 (381.28)	13.26	885	248	227
Exports	-69.52 (137.62)	23.06	260	109	135
Investment	2.42 (22.57)	25.50	707	381	326
Employment	-1.72 (3.11)	18.05	885	365	291
		B: 2000			
Turnover	-792.25 (420.74)	14.46	919	293	244
Exports	-66.03 (166.28)	12.57	234	47	71
Investment	-18.87 (23.33)	14.87	701	240	187
Employment	-2.85 (3.58)	16.55	919	348	278
		C: 2001			
Turnover	81.81 (447.33)	13.73	932	265	264
Exports	-75.62 (110.28)	15.33	226	55	79
Investment	-6.67 (18.90)	15.85	642	206	214
Employment	-3.75 (3.50)	9.36	932	195	190
		D: 2002			
Turnover	97.90 (378.42)	22.94	956	450	438
Exports	41.45 (157.28)	16.23	226	64	80
Investment	-9.11 (14.20)	19.69	632	241	278
Employment	1.99 (2.35)	19.46	955	385	386
		E: 2003			
Turnover	332.94 (445.92)	12.78	577	138	192
Exports	-260.35 (415.11)	15.42	130	35	54
Investment	-7.78 (19.83)	13.00	433	94	148
Employment	-0.12 (2.49)	17.09	577	176	272
		F: 2004			
Turnover	367.21 (295.79)	19.59	528	164	295
Exports	180.58 (253.08)	12.80	108	19	41
Investment	-17.15 (16.48)	11.67	357	72	118
Employment	1.71 (2.54)	14.06	528	126	208
		G: 2005			
Turnover	825.30* (382.53)	14.23	498	107	189
Exports	402.82 (228.55)	16.12	109	24	45
Investment	-10.67 (10.59)	13.90	330	76	109
Employment	-0.31 (2.14)	22.87	498	146	321

Notes: * indicates significance at the 5 percent level. Standard errors are shown in parentheses. Turnover, investment and exports are denoted in EUR 1,000. The number of observations refer to the ± 25 MWh region around the threshold c. The bandwidth is selected based on the procedure in The bandwidth is selected based on the procedure in Imbens and Kalyanaraman (2012). Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey, own calculations.

4.5.7 Updated decision making

In this section, we relax the assumption that firms make accurate production decisions based on a marginal cost function that is constructed at the beginning of each year given the available information. Instead of assuming that firms correctly predict their electricity use for the coming year, we assume firms to frequently update (e.g. quarterly or monthly) their production decisions based on newly available information. In this case, a firm would change its production decision as soon as it exceeds the threshold for the electricity tax reduction.

Assuming that firms update their production decisions modifies our identification strategy. The discontinuity in the marginal tax rate does not translate anymore into a discontinuity in the conditional expectation of the outcome variable given electricity use. Instead, updated decision making causes a change in the slope of the functional relationship between the outcome variable and electricity use. We expect a potential change in the slope to be particularly strong for value added, because marginal production costs decrease once a firm exceeds the energy use threshold for the reduced electricity tax rate. To see this, suppose that a firm uses 50 MWh electricity in 2003. Given our assumption, the firm adjusts its production decision as soon as its electricity use exceeds the threshold of 25 MWh. As a consequence, costs related to the use of energy decrease for every unit that is additionally used. Examining the data on annual base, the discontinuity in the marginal tax rate then translates into a kink in the conditional expectation of value added given electricity use.

We estimate an alternative specification of our model taking updated decision making into account. Instead of investigating whether there is a discontinuity in the conditional expectation of value added given electricity use, we examine whether the slope of this functional relationship changes in the neighborhood of the thresholds. In particular, we take the difference of the slope parameters $\hat{\beta}_+$ and $\hat{\beta}_-$ as an estimate for the local average treatment effect (LATE). These parameters are obtained by estimating two weighted local linear regression as described in Section 4.3.3.

In Table 4.8, we show the estimates and robust standard errors for the kink LATE estimates for the bandwidths 5, 10, 15, 20, and 25 MWh, since standard IK bandwidths do not apply for this model specification (Card, Lee, Pei, and Weber, 2015).[16] The estimated parameters reflect how much more (or less) value added (in 1,000 euros) is generated by the treatment group when using one additional unit of electricity (in MWh).

We do not find any statistically significant effects of the electricity tax reduction on value added under updated decision making. The standard errors decrease with increasing bandwidth, while the estimates tend to converge to zero. This can be partly explained by the number of observations, that lies between 50 and 100 for the 5 MWh bandwidth and is thus rather low. For the 25 MWh bandwidth, the number of observations lies between 150 and 700 depending on the year (and threshold) under investigation. Also following this approach, we do not find a consistent pattern of negative or positive effects.

Table 4.8: Kink LATE estimates for value added

Value added	Bandwidth (in MWh)				
Year	5	10	15	20	25
1999	525.25 (644.58)	297.77 (206.46)	-25.92 (104.14)	-42.54 (68.78)	-21.993 (50.789)
2000	-622.74 (583.59)	-78.54 (218.79)	-7.22 (118.42)	-2.20 (73.27)	-0.06 (52.07)
2001	450.37 (483.67)	124.35 (208.35)	-40.52 (123.27)	-25.34 (78.65)	29.68 (57.99)
2002	929.33 (1146.80)	223.81 (253.71)	77.205 (165.03)	56.836 (112.71)	75.873 (83.89)
2003	48.45 (477.20)	-8.23 (167.27)	12.27 (91.66)	24.54 (64.63)	54.65 (52.31)
2004	873.93 (407.65)	47.81 (187.29)	-35.77 (118.74)	-56.29 (79.56)	-20.90 (64.10)
2005	-378.62 (918.29)	-143.88 (282.37)	-104.13 (160.25)	-53.378 (105.37)	-14.003 (75.35)

Notes: * indicates significance at the 5 percent level. Standard errors are shown in parentheses. Value added is denoted in EUR 1,000. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey own calculations.

[16] With regard to this specification, we expect the strongest effect for value added. For this reason, we focus on the impact of the change in marginal tax rates on value added.

4.6 Discussion

In this section, we discuss several factors that may have influenced our findings. Thereby we also draw attention to related and future research. First, we discuss the statistical power of our analysis. Then, we assess the likelihood and implications of a possible violation of the Stable Unit Treatment Value Assumption (SUTVA). Finally, we debate how our local results may relate to a wider set of firms.

Several factors influence the power of a statistical analysis, i.e., the correct rejection of the null hypothesis of no effects, when it is false. While some factors suggest that the power of our analysis is high, others suggest the opposite, with neither side clearly dominating.

First, we discuss the magnitude of the effect. If the size of effects is small, statistical power tends to be low. In our case, the electricity tax strongly changes the price of electricity. During the period under investigation, it increases the pre-tax electricity price by 15 to 27 percent on average as shown in Section 4.2. This is a large change suggesting an effect of significant magnitude. The change in electricity price is also large in comparison to the Climate Change Levy (CCL) in the United Kingdom, for which Martin, de Preux, and Wagner (2014) did neither find any negative effects on economic outcomes. At its introduction in 2001, the CCL amounted to GPB 4.35 per MWh, or 7 to 11 percent of the pre-tax electricity price (Eurostat, 2014, own calculations).

Electricity is only one of many inputs to production. Even a strong change in the electricity price may thereby have a limited impact on firms suggesting a small magnitude of the effect. For this reason we estimate our model for the steel sector in Section 4.5.6, which is electricity intensive and thereby particularly exposed to changes in electricity prices. Nevertheless, we do not find evidence for a causal effect of the electricity tax reduction in the steel sector.

Second, our data is fairly heterogeneous. This leads to a risk of not rejecting the null hypothesis although the null hypothesis is false for at least some firms. To account for such a possibility, we analyzed different homogenous subpopulations in Section 4.5.6. While the precision of the estimates tends to increase, we do not find robust evidence of significant effects in the subpopulations.

Third, we draw our attention to sample size and measurement error. Low sample size and high measurement error would suggest low statistical power. Except for value added, our data are based on censuses. While also data collected by mandatory census can exhibit measurement error, we are provided with data for all firms with more than 20 employees, i.e. sampling errors do not affect our results. In addition, the number of observations is typically large.

We have not yet discussed to what extent effects on the treated firms may induce additional effects on untreated firms. If such effects would occur, SUTVA would be violated. In the following paragraphs we discuss a likely violation of SUTVA, what its effect would be, and if we can find any evidence for such a violation.

The interaction of treated and untreated firms in common markets may violate SUTVA. Let us assume that there was a positive direct effect of the reduced tax rate on turnover for a treated firm, as marginal production costs have decreased compared to the level of the full tax rate, and lower production costs enable higher production levels. If this treated firm is in competition with another untreated firm in the same market, the treatment may have spill-over effects to the untreated firm. In particular, the treated firm may gain additional market share by lowering the product price to a level where the untreated firm that has higher marginal costs cannot compete. In such a situation, the positive spill-over effect would add to the positive direct effect of the tax reduction.

While we are not able to distinguish for a single year what part of the total effect consists of the direct effect of the reduced tax rate or the spill-over effects from being able to gain market share through altering prices, we can assess whether hypothesized effects are particularly strong for the year when the treatment was strongest. Going back to Table 4.3 in Section 4.5.3 we do not observe particularly strong effects for the year 2002 when the difference between the full and the reduced tax rate was highest, in particular when dividing total effects by the size of the tax reduction. Furthermore, effect signs are mostly negative, which is not in line with a positive spill-over effect due to reduced marginal costs. In addition estimates are statistically insignificant except for a negative coefficient for investments. Taken together, we do not observe strong evidence that SUTVA is violated due to spill-over effects.

Last but not least, we debate how our local results may relate to a wider set of firms. Looking back at Table 4.2 and Figure 4.2 in Section 4.4.3, the analyzed firms fall within the lower quintile of energy use. While small, energy-intensive firms as well as larger, less energy-intensive firms are covered by our analysis, large energy-intensive firms are hardly covered. This raises the question whether our results would also apply to large, energy-intensive firms. It is not unlikely that larger electricity intensive firms differ from the firms under investigation, for example, with respect to own electricity generation. The best way forward may be to look out for similar experiments in tax rates or levies that do apply to larger firms.

A related question is in how far our results are relevant for policy, given that we assessed the effects of a tax reduction for relatively small firms in terms of electricity use. It should be noted that the tax reduction was granted precisely in order to mitigate any negative impacts on firm's performance and particularly exports. Given that we do not find any positive effects of a reduced tax rate on firm's performance, or in other words, any negative effects of higher electricity taxes, this puts doubts on the necessity of the tax reduction for domestic economic reasons. While we cannot rule out that large, energy-intensive firms may be affected differently than smaller firms by the electricity tax, we can say at least that the tax reduction is not well targeted for its purpose. Tax revenues are forgone by providing relief to firms that are not found to be vulnerable to higher electricity taxes.

4.7 Conclusion

This paper analyzed the causal impacts of the German electricity tax on the economic performance of firms in the manufacturing sector. The tax was implemented in 1999 and firms with electricity use above a certain threshold were eligible for a reduced electricity tax rate. We evaluated the effects of the reduced marginal electricity tax rate on five variables of economic performance, namely, turnover, exports, investment, employment, and value added with a regression discontinuity analysis. No robust positive or negative impact of the reduced marginal electricity tax rate was found. Hence, our results indicate that firms forced to pay the full electricity tax rate did not suffer from deterioration in their economic performance.

Our findings suggest that the reduced electricity tax rate may not be needed for securing the competitiveness of firms in the manufacturing sector. Firms that had to pay the higher electricity tax did not perform worse than firms that only had to pay the reduced electricity tax rate. It can thus be expected that firms that have had to pay the reduced electricity tax would also adjust smoothly, if the tax reduction was removed. If there are doubts about the ability of firms with substantially higher electricity-use than investigated to adjust to the higher electricity tax rate, the electricity tax rate they pay could be increased stepwise. The threshold for eligibility of the reduced tax rate could be increased, accompanied by a causal evaluation of its impacts on the economic performance of firms. Removing the reduced tax rate would raise revenues for the government that could be used to decrease more distorting taxes, to consolidate budgets, or to finance new investments.

4.8 Acknowledgements

We thank Andreas Löschel, Kathrine von Graevenitz, Kurt van Dender, Maximilian Auffhammer, Ralf Martin, Ulrich Wagner, Meredith Fowlie, Nils Axel Braathen, Katrin Sommerfeld, Nico Pestel, Tomasz Kozluk, Jonathan Goyette, Andreas Gerster, Beat Hintermann, Stefan Lamp, and Stefan Vögele for suggestions and insightful comments. We gratefully acknowledge the Research Data Centre (FDZ) of the Federal Statistical Office and the Statistical Offices of the German Länder for granting us access to the AFiD data and for the use of their research facilities, in particular, Michael Rößner for his advice and technical support regarding the estimation of our empirical model. The paper has also benefited from discussions with participants of the ASSA Meeting 2015, the 6^{th} AWEEE 2014, the 37^{th} IAEE International Conference 2014, and the 19^{th} SMYE 2014. We also thank the participants of internal seminars and meetings at OECD, ZEW, as well as UC Berkeley and colleagues at OECD and ZEW for informal discussions. Benjamin Johannes Lutz thanks the Helmholtz Association for the financial support through the Helmholtz Alliance ENERGY-TRANS. The views expressed in this paper are those of the authors and do not necessarily represent those of the OECD, ZEW or the FDZ.

4.9 Appendices

The following appendices provide additional information on electricity intensity, calculated as the ratio of electricity use and turnover, the effects of kernel and bandwidth choices, lagged treatment effects as well as descriptive statistics for the subpopulation analysis.

Appendix 4.A Additional information on electricity intensity

Below, we depict the relationship between electricity intensity and total electricity use showing scatter plots of non-overlapping binned local means and second order global polynomial functions of electricity intensity in Figures 4.9 and 4.10. The local means are computed for 1 MWh bandwidths in the area of 25 - 75 MWh, the $c \pm 25$ MWh neighborhood surrounding the prevailing threshold. Given that electricity intensity is computed as the ratio of electricity use and turnover, these relationships do not have any causal interpretation.

In order to depict the relationship between electricity intensity and total electricity use, we show scatter plots of non-overlapping binned local means and second order global polynomial functions of the variable electricity intensity in Figures 4.9 and 4.10. The local means are computed for 1 MWh bandwidths in the area of 25 - 75 MWh, the $c \pm 25$ MWh neighborhood surrounding the prevailing threshold.

Figure 4.9: Outcomes in year 1995

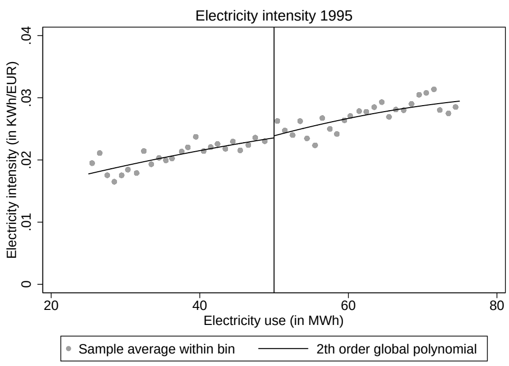

Notes: Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

Figure 4.10: Electricity use and intensity

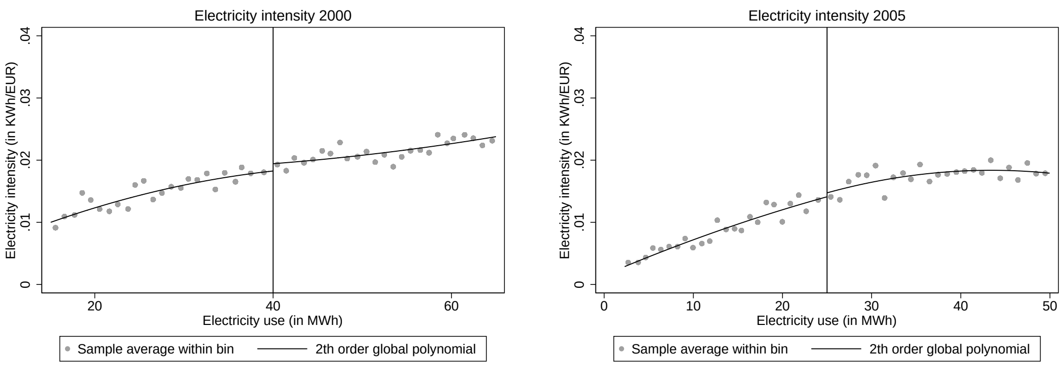

Notes: Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units and AFiD-Module Use of Energy, own calculations.

Appendix 4.B The effect of alternative kernel choices

In this section, we provide evidence for the robustness of our findings with regard to the kernel choice. We show local average treatment effects using uniform (Table 4.9) and Epanechnikov kernel functions (Table 4.10) for the local linear regressions. The alternative kernel choice does not change qualitative results.

Table 4.9: Uniform kernel function

Outcome variable	Estimator	Bandwidth	Number of observations		
			c ± 25 MWh	Control group	Treatment group
		A: 1999			
Turnover	139.38 (186.33)	17.43	5,289	1,953	1,807
Exports	-85.83 (122.87)	10.96	2,330	489	554
Investment	-5.75 (15.91)	12.11	3,873	1,342	1,277
Employment	-0.42 (1.06)	17.82	5,289	2,002	1,835
Value added	-125.60 (213.34)	15.41	1,452	488	456
		B: 2000			
Turnover	-157.45 (193.11)	16.27	5,017	1,743	1,615
Exports	-32.13 (121.35)	12.90	2,137	539	562
Investment	-1.55 (12.60)	12.87	3,691	1,001	950
Employment	-0.05 (1.24)	13.72	5,014	1,434	1,338
Value added	-93.90 (219.98)	21.07	1,301	387	382
		C: 2001			
Turnover	443.39* (286.44)	12.75	4,862	1,242	1,347
Exports	-145.09 (133.01)	12.29	2,041	478	540
Investment	7.56 (10.74)	12.30	3,338	842	898
Employment	0.50 (1.03)	20.10	4,859	1,829	2,077
Value added	148.56 (235.13)	14.89	1,119	312	369
		D: 2002			
Turnover	-520.72* (265.37)	9.51	5,072	937	1,021
Exports	-97.99 (114.00)	13.93	2,114	536	630
Investment	-20.82 (10.78)	16.13	3,360	970	1,197
Employment	0.46 (1.26)	14.73	5,063	1,390	1,584
Value added	-297.42 (303.96)	16.82	985	310	356
		E: 2003			
Turnover	-205.82 (227.5)	12.74	3,052	730	964
Exports	-237.33 (215.91)	9.52	1,290	230	290
Investment	-3.69 (9.22)	12.89	2,175	502	708
Employment	-0.83 (1.41)	13.89	3,052	780	1,066
Value added	-142.88 (206.11)	11.32	851	190	239
		F: 2004			
Turnover	278.38 (233.69)	10.42	2,779	527	776
Exports	-136.71 (160.78)	11.10	1,138	198	354
Investment	-1.66 (9.14)	13.20	1,979	452	706
Employment	0.50 (1.13)	13.19	2,778	621	994
Value added	102.98 (237.72)	11.18	704	138	226
		G: 2005			
Turnover	-172.18 (275.96)	9.73	2,654	450	680
Exports	306.72 (181.43)	12.28	1,068	211	314
Investment	17.17* (8.39)	13.30	1,856	412	629
Employment	0.52 (1.41)	9.29	2,654	431	649
Value added	-46.01 (215.49)	17.48	621	157	303

Notes: * indicates significance at the 5 percent level. Standard errors are shown in parentheses. Turnover, investment, and exports are denoted in EUR 1,000. The number of observations refer to the ± 25 MWh region around the threshold c. The bandwidth is selected based on the procedure in Imbens and Kalyanaraman (2012). Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey own calculations.

Table 4.10: Epanechnikov kernel function

Outcome variable	Estimator	Bandwidth	Number of observations		
			c ± 25 MWh	Control group	Treatment group
		A: 1999			
Turnover	99.96 (174.23)	22.00	5,289	2,463	2,192
Exports	-15.30 (112.03)	14.41	2,330	680	707
Investment	-8.84 (11.45)	20.79	3,873	1,682	1,557
Employment	-0.45 (1.01)	21.63	5,289	2,431	2,156
Value added	-73.72 (202.47)	19.42	1,452	598	556
		B: 2000			
Turnover	-162.95 (182.63)	20.43	5,017	2,121	2,076
Exports	-29.86 (110.36)	16.80	2,137	706	752
Investment	-3.48 (11.44)	17.35	3,691	1,372	1,280
Employment	-0.06 (1.17)	17.58	5,014	1,867	1,731
Value added	1.98 (204.78)	18.37	1,301	499	477
		C: 2001			
Turnover	436.22* (218.88)	15.37	4,862	1,473	1,651
Exports	-160.35 (124.85)	16.02	2,041	591	742
Investment	9.58 (10.10)	15.59	3,338	1,003	1,175
Employment	0.64 (0.97)	24.41	4,859	2,300	2,485
Value added	165.92 (215.22)	18.61	1,119	383	460
		D: 2002			
Turnover	-394.9 (239.91)	13.20	5,072	1,259	1,412
Exports	-70.82 (111.38)	16.56	2,114	612	757
Investment	-19.48 (10.283)	19.33	3,360	1,264	1,419
Employment	0.09 (1.14)	18.49	5,063	1,925	1,973
Value added	-446.43 (301.25)	20.60	985	354	469
		E: 2003			
Turnover	-161.33 (222.86)	15.13	3,052	861	1,193
Exports	-234.40 (152.83)	11.90	1,290	265	376
Investment	-4.24 (8.56)	16.69	2,175	618	979
Employment	-0.63 (1.33)	17.35	3,052	917	1,399
Value added	-182.57 (186.50)	14.75	851	232	320
		F: 2004			
Turnover	235.80 (219.85)	13.03	2,779	620	978
Exports	-65.13 (156.85)	13.36	1,138	222	431
Investment	-3.53 (8.91)	16.09	1,979	527	882
Employment	0.54 (1.06)	16.70	2,778	743	1,276
Value added	90.30 (229.36)	14.59	704	164	296
		G: 2005			
Turnover	-125.48 (270.40)	11.33	2,654	502	781
Exports	328.80* (166.92)	15.86	1,068	256	438
Investment	14.96 (7.97)	16.06	1,856	479	798
Employment	0.50 (1.34)	11.52	2,654	505	790
Value added	-18.37 (211.83)	21.35	621	175	373

Notes: * indicates significance at the 5 percent level. Standard errors are shown in parentheses. Turnover, investment, and exports are denoted in EUR 1,000. The number of observations refer to the ± 25 MWh region around the threshold c. The bandwidth is selected based on the procedure in Imbens and Kalyanaraman (2012). Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey own calculations.

Appendix 4.C The effects of bandwidth choice

In order to examine the sensitivity of results to different bandwidth choices, we estimate the local average treatment effect for 5, 10, 15, 20, and 25 MWh bandwidths. Figure 4.11 shows the local average treatment effects of the tax reduction on turnover in 2001 and investment in 2002 as a function of bandwidth choice. Table 4.11 shows the results for the treatment years 1999 - 2005.

Figure 4.11: The effect of bandwidth choice on point estimates and confidence intervals

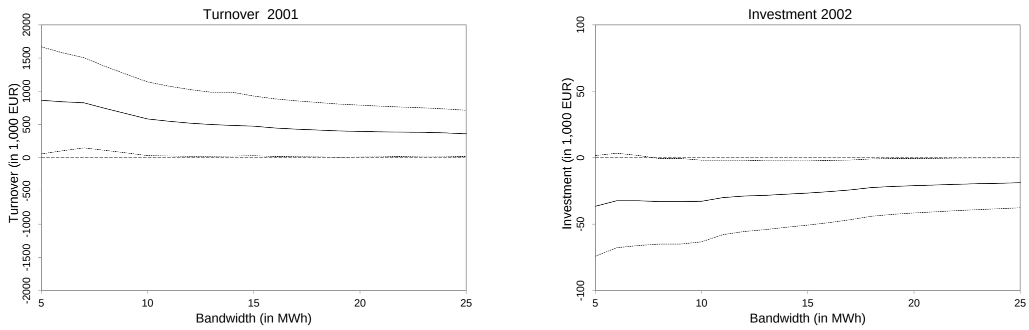

Notes: The solid black line in each graph denotes point estimates and the dashed lines are corresponding 95 percent confidence intervals. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey, own calculations.

Table 4.11: LATE estimates for 5, 10, 15, 20, and 25 MWh bandwidths

Outcome variable	Bandwidth (in MWh)				
	5	10	15	20	25
			A: 1999		
Turnover	117.48 (317.93)	239.87 (244.65)	110.45 (207.22)	119.26 (185.22)	96.74 (167.17)
Exports	-42.64 (205.45)	-70.05 (137.65)	-18.05 (112.55)	26.74 (97.59)	24.70 (86.83)
Investment	-25.07 (24.10)	-15.78 (16.53)	-10.43 (13.66)	-9.00 (12.02)	-10.84 (10.73)
Employment	0.96 (1.74)	-0.49 (1.37)	-0.68 (1.18)	-0.50 (1.06)	-0.31 (0.96)
Value added	-393.99 (366.91)	-127.31 (272.50)	-101.02 (231.67)	-82.34 (205.11)	-63.02 (183.16)
			B: 2000		
Turnover	-450.60 (407.13)	-261.96 (276.13)	-189.04 (221.20)	-168.92 (191.06)	-154.24 (171.41)
Exports	-278.31 (219.16)	-215.61 (151.92)	-95.03 (122.01)	-17.12 (104.62)	-0.34 (92.90)
Investment	26.53 (27.68)	5.31 (17.51)	-0.36 (13.47)	-1.07 (11.24)	1.27 (9.87)
Employment	-1.68 (2.71)	-1.59 (1.75)	-0.65 (1.35)	-0.07 (1.15)	0.10 (1.02)
Value added	-649.01 (370.99)	-377.52 (267.74)	-161.32 (225.62)	-20.71 (201.03)	44.75 (183.63)
			C: 2001		
Turnover	862.92* (410.84)	585.95* (282.79)	479.47* (229.00)	401.60* (198.85)	365.53* (177.81)
Exports	-447.05* (221.06)	-196.72 (159.24)	-163.59 (132.39)	-203.16 (115.62)	-198.11 (103.18)
Investment	-8.46 (16.92)	5.71 (12.526)	8.26 (10.46)	8.55 (9.28)	8.62 (8.41)
Employment	0.31 (2.10)	0.25 (1.51)	0.51 (1.25)	0.41 (1.10)	0.59 (0.98)
Value added	512.99 (361.72)	410.13 (261.56)	251.77 (231.15)	187.23 (209.65)	140.40 (192.57)
			D: 2002		
Turnover	-582.48 (369.92)	-458.00 (278.24)	-347.30 (231.72)	-109.13 (199.41)	-20.128 (182.54)
Exports	144.04 (247.47)	6.66 (152.31)	-46.47 (121.53)	-42.79 (102.78)	-28.10 (93.50)
Investment	-36.21 (19.31)	-31.75* (14.97)	-26.53* (12.33)	-21.04* (10.47)	-18.90* (9.57)
Employment	1.55 (2.18)	-0.49 (1.62)	-0.45 (1.34)	0.126 (1.14)	0.20 (1.03)
Value added	-972.51* (495.41)	-874.32* (444.76)	-746.64* (367.99)	-528.14 (313.00)	-481.94 (285.39)
			E: 2003		
Turnover	33.23 (401.90)	23.86 (286.46)	-125.67 (231.83)	-186.62 (201.78)	-215.76 (185.30)
Exports	-436.06 (266.66)	-265.37 (182.88)	-181.41 (139.74)	-68.66 (117.95)	-32.77 (108.96)
Investment	12.59 (14.15)	-0.67 (11.10)	-2.31 (9.28)	-4.73 (8.15)	-4.02 (7.58)
Employment	1.92 (2.66)	0.80 (1.87)	-0.26 (1.50)	-0.53 (1.30)	-0.55 (1.19)
Value added	-80.35 (307.41)	-173.59 (226.77)	-175.64 (189.31)	-202.80 (166.43)	-185.66 (53.37)
			F: 2004		
Turnover	711.13* (322.74)	332.62 (250.54)	236.29 (210.94)	125.90 (186.67)	63.04 (172.21)
Exports	15.37 (269.11)	-87.37 (193.97)	-39.11 (154.74)	39.89 (132.73)	98.11 (121.01)
Investment	4.19 (17.16)	-4.57 (12.48)	-2.48 (9.83)	-5.42 (8.40)	-6.07 (7.69)
Employment	5.03* (1.77)	2.72* (1.36)	1.15 (1.13)	0.64 (1.01)	0.63 (0.94)
Value added	137.17 (392.29)	135.98 (292.62)	93.51 (237.09)	52.87 (204.60)	56.77 (186.15)
			G: 2005		
Turnover	137.01 (378.05)	-79.77 (290.82)	-149.71 (244.83)	-216.97 (216.96)	-309.6 (198.33)
Exports	448.81 (267.83)	362.2 (203.21)	323.35 (174.51)	345.12* (154.47)	327.60* (140.93)
Investment	17.57 (12.47)	10.28 (9.83)	14.21 (8.41)	13.70 (7.48)	10.88 (6.86)
Employment	3.02 (1.95)	1.10 (1.45)	0.31 (1.23)	0.39 (1.10)	0.52 (1.02)
Value added	-569.45 (552.83)	-189.10 (337.20)	-56.01 (262.19)	-35.54 (228.33)	-42.74 (208.33)

Notes: * indicates significance at the 5 percent level. Standard errors are shown in parentheses. Turnover, exports, investments, and value added are denoted in EUR 1,000. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey own calculations.

Appendix 4.D Dynamic local average treatment effects

Table 4.12 shows lagged local average treatment effects. The estimators indicate that the results of our analysis are robust with respect to potential adjustment processes that might lead to delayed effects of the tax reduction.

Table 4.12: Dynamic local average treatment effects

Outcome variable	Estimator	Bandwidth	Number of observations		
			c ± 25 MWh	Control group	Treatment group
A: Effect of the discontinuity in 1999 on outcome in 2000					
Turnover	7.47 (217.70)	21.07	4,672	2,053	1,893
Exports	-73.96 (114.51)	22.39	2,107	919	997
Investment	8.75 (12.92)	21.41	3,575	1,595	1,477
Employment	-1.35 (1.12)	21.09	4,665	2,052	1,896
Value added	-133.27 (208.29)	19.05	1,375	559	527
B: Effect of the discontinuity in 2000 on outcome in 2001					
Turnover	68.76 (189.15)	30.97	4,403	2,179	2,224
Exports	-156.23 (146.81)	16.98	1,900	633	663
Investment	7.24 (12.40)	19.13	3,215	1,270	1,245
Employment	0.47 (1.11)	28.057	4,403	2,185	2,218
Value added	-137.59 (250.69)	14.55	1,151	333	344
C: Effect of the discontinuity in 2001 on outcome in 2002					
Turnover	524.44* (259.06)	16.93	4,148	1,312	1,578
Exports	-66.78 (137.73)	19.96	1,749	580	811
Investment	2.52 (10.44)	21.30	2,891	1,223	1,369
Employment	0.30 (1.18)	22.13	4,148	1,794	1,997
Value added	420.40 (308.58)	15.02	1,020	284	337
D: Effect of the discontinuity in 2002 on outcome in 2003					
Turnover	-240.08 (235.97)	17.929	4,255	1,574	1,631
Exports	29.96 (143.29)	18.77	1,862	672	774
Investment	-3.41 (9.29)	19.86	2,983	1,145	1,324
Employment	-0.28 (1.33)	19.53	4,255	1,649	1,830
Value added	-159.47 (337.57)	19.67	1,149	452	497
E: Effect of the discontinuity in 2003 on outcome in 2004					
Turnover	-235.37 (224.35)	16.97	2,842	853	1,271
Exports	-259.18 (217.54)	10.56	1,195	234	315
Investment	3.55 (10.53)	19.23	1,986	518	759
Employment	-2.26 (1.34)	18.80	2,842	891	1,421
Value added	-320.15 (241.86)	15.02	780	213	304
F: Effect of the discontinuity in 2004 on outcome in 2005					
Turnover	440.04 (279.36)	12.60	2,572	564	864
Exports	164.28 (173.88)	20.92	1,067	287	643
Investment	15.18 (12.22)	13.52	1,749	410	621
Employment	2.15 (1.19)	16.25	2,571	682	1,144
Value added	179.16 (311.78)	11.08	645	125	198
G: Effect of the discontinuity in 2005 on outcome in 2006					
Turnover	-274.6 (279.77)	14.07	2,393	546	868
Exports	383.38 (208.65)	14.79	993	220	358
Investment	9.41 (12.10)	14.33	1,729	396	638
Employment	1.12 (1.47)	13.74	2,392	527	853
Value added	-9.16 (267.04)	17.31	577	149	279

Notes: * indicates significance at the 5 percent level. Standard errors are shown in parentheses. Turnover, investment, and exports are denoted in EUR 1,000. The number of observations refer to the ± 25 MWh region around the threshold c. The bandwidth is selected based on the procedure in Imbens and Kalyanaraman (2012). Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey own calculations.

Appendix 4.E Subpopulation analysis: descriptive statistics

The subpopulation manufacture of machinery, electronic devices, and vehicles covers the industries 29 - 35 according to the ISIC Rev. 3.1 classification. In Table 4.13, we present the descriptive statistics of the assignment variable electricity use, electricity intensity, and the outcome variables considered in Section 4.5.6. Firms that manufacture machinery, electronic devices, and vehicles show on average higher turnovers and exports and are less electricity intensive in comparison to the full population. In terms of turnover and number of employees, the subpopulation is less heterogeneous.

The subpopulation manufacture of basic metal and fabricated metal products comprises the industries 27 and 28 according to the ISIC Rev. 3.1 classification. The descriptive statistics in Table 4.14 show, that firms of this subpopulation produce less output in terms of turnover and export less than the average firm of the full population. Furthermore, when comparing average electricity intensities, we see that this industry is on average more energy intensive than the full population. With regard to turnover, number of employees, and electricity intensity, the subpopulation is less heterogenous.

Table 4.13: Descriptive statistics: manufacture of machinery, electronic devices, and vehicles.

	Mean	St. dev.	P10	P 50	P90	N
	A: 1999					
Electricity use (in MWh)	972.99	2,716.51	32.8	206.78	2188.13	10,758
Turnover (in EUR 1,000)	14,400.29	24,750.78	1,561.58	5,644.63	35,646.88	10,769
Exports (in EUR 1,000)	2,622.11	4,407.23	0	606.87	12,316.20	10,769
Investment (in EUR 1,000)	538.261	1,378.32	0	120.45	1,326.61	9,429
Employment	126.72	184.55	25.08	58.75	296.42	10,769
Electricity intensity (in EUR per KWh)	0.0615	0.0764	0.0087	0.0379	0.1385	10,758
	B: 2000					
Electricity use (in MWh)	1,028.71	2,978.79	33.30	216.10	2,283.79	11,319
Turnover (in EUR 1,000)	16,973.42	29,564.06	1,745.99	6,376.33	42,803.77	11,324
Exports (in EUR 1,000)	6,104.53	14,432.49	0	877.08	16,644.36	11,324
Investment (in EUR 1,000)	591.60	1,381.61	0.21	137.20	1,483.91	10,914
Employment	117.48	161.77	24.83	56.21	277.58	11,324
Electricity intensity (in EUR per KWh)	0.0573	0.0713	0.0077	0.0345	0.1290	11,319
	C: 2005					
Electricity use (in MWh)	1,189.87	3,438.61	45.44	262.49	2,603.56	11,334
Turnover (in EUR 1,000)	18,332.25	32,909.89	1,671.67	6,660.25	45,592.08	11,750
Exports (in EUR 1,000)	7,581.08	17,913.27	0	1,209.86	20,499.69	11,750
Investment (in EUR 1,000)	465.87	1169.16	0	94.88	1,143.12	11,287
Employment	112.81	155.22	24.67	55.00	263.42	11,750
Electricity intensity (in EUR per KWh)	0.0653	0.0847	0.0104	0.0383	0.1488	11,259

Notes: Turnover, investment, and exports are denoted in EUR 1,000. Electricity use relates to the taxable electricity use in MWh (not including self-generated electricity). Electricity intensity is denoted by electricity use divided by turnover. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey, own calculations.

Table 4.14: Descriptive statistics: manufacture of basic metal and fabricated metal products.

	Mean	St. dev.	P10	P 50	P90	N
	A: 1999					
Electricity use (in MWh)	1,541.63	4,020.85	35.05	317.32	3,512.14	6,477
Turnover (in EUR 1,000)	11,596.03	20,565.58	1,640.40	4,873.48	26,272.04	6,482
Exports (in EUR 1,000)	2,122.57	6,644.29	0	68.06	5,098.35	6,482
Investment (in EUR 1,000)	485.21	1,131.52	0	121.79	1,191.25	5,810
Employment	99.90	143.96	24.58	51.92	215	6,482
Electricity intensity (in EUR per KWh)	0.1165	0.1417	0.0103	0.0687	0.2857	6,477
	B: 2000					
Electricity use (in MWh)	1,686.35	4,518.33	36.00	347.90	3844.52	6,986
Turnover (in EUR 1,000)	12,415.21	22,695.61	1,730.20	5,018.85	28,806.22	6,994
Exports (in EUR 1,000)	2,843.71	9,110.51	0	98.87	6,826.69	6,994
Investment (in EUR 1,000)	559.73	1,343.42	0	125.13	1,341.16	6,752
Employment	91.65	127.76	24.08	48	197.75	6,994
Electricity intensity (in EUR per KWh)	0.1167	0.1397	0.0105	0.0709	0.2794	6,986
	C: 2005					
Electricity use (in MWh)	2,090.48	5,228.09	62.16	480	4,826.61	6,783
Turnover (in EUR 1,000)	13,319.77	24,694.88	1,709.18	5,102.57	31,850.46	6,963
Exports (in EUR 1,000)	3,680.77	10,802.40	0	190.61	9,462.13	6,963
Investment (in EUR 1,000)	428.52	1,060.49	0	83.61	1,078.46	6,668
Employment	88.19	121.65	24	47.17	189.67	6,963
Electricity intensity (in EUR per KWh)	0.1428	0.1614	0.0180	0.0922	0.3239	6,737

Notes: Turnover, investment, and exports are denoted in EUR 1,000. Electricity use relates to the taxable electricity use in MWh (not including self-generated electricity). Electricity intensity is denoted by electricity use divided by turnover. Source: Research Data Centres of the Statistical Offices Germany (2014): Official Firm Data for Germany (AFiD) - AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey, own calculations.

Chapter 5

Nonlinearity in cap-and-trade systems: the EUA price and its fundamentals

5.1 Introduction

The EUA price dynamics and its driving factors have been of great importance for practitioners, politicians, and scientists since the introduction of the European Union Emissions Trading Scheme (EU ETS) in 2005. The reasons for the interest are manifold. First, carbon prices introduce an additional cost component affecting day-to-day and long-term operations of regulated installations. Understanding this cost component is a key strategic element for many regulated installations to achieve long-term cost efficiency. Second, the scheme is a market-based policy instrument. Its success heavily depends on its ability to generate correct price signals that fully account for the underlying fundamentals. Thus, the relation between the EUA price and its driving factors is crucial for the understanding of the effectiveness of the scheme.

We argue that the varying relation between the EUA price and its fundamentals is a consequence of the design of the EU ETS. In cap-and-trade systems, as in the case of the EU ETS, the regulatory authority determines the total number of allowances for a certain period of time. In other words, the aggregated supply of allowances is fixed and therefore inelastic. In contrast, the demand varies due to various shocks, for example positive and negative shocks to the macroeconomic activity. Such shocks shift the production of goods to higher or lower levels, which increases or decreases the aggregated level of emissions and, thus, the demand for allowances. As a result, market participants adjust their expectations about the overall stringency of the scheme. We hypothesize, that this situation translates into a higher volatility and a varying relation between the EUA price and its fundamentals.

The recent literature provides empirical evidence on structural changes in the data generating process of the EUA prices. Alberola, Chevallier, and Chèze (2008), Chevallier (2009), Keppler and Mansanet-Bataller (2010), and Hintermann (2010) devote their research to the detection of price determinants affecting the European carbon market. In particular, they quantify the linear impact of fundamentals such as commodity prices, weather conditions, and economic fluctuations on the EUA price. To account for potentially time-changing influences of the fundamentals, they conduct an analysis over different subsamples. In doing so, they assume that the timing of the structural breaks is known. The authors trace these structural changes back to different factors: Alberola, Chevallier, and Chèze (2008) refer to the information disclosure on the actual emissions in 2006 as a reason for structural changes, whereas Chevallier (2009) sees the aftermath of the financial crisis as a factor causing breaks. These potential sources for breaks seem to have one characteristic in common: They alter the expectations about the overall demand of allowances during the prevailing compliance period and, thus, affect the expectations about the overall stringency of the EU ETS. Therefore, these changes are inherent in the cap-and-trade system and should be endogenized.

In contrast to these earlier studies, we therefore do not assume the changes in the regimes to be deterministic. Instead, we consider a Markov regime-switching model and simultaneously allow for both: time-variation in the effects of the fundamentals as well as in the volatility of the EUA prices, as changes in expectations about the overall stringency of the EU ETS are likely to be associated with periods of higher uncertainty and, thus, with higher volatility in the EUA prices. Moreover, we focus on the short-term fluctuations and consider the entire Kyoto commitment period, i.e. Phase II, which ranges from 1 January 2008 until 31 December 2012.

In doing so, we contribute to the more recent literature employing nonlinear models to examine the relationship between the EUA price and its determinants. Recent findings of Chevallier (2011a), Chevallier (2011b) and Peri and Baldi (2011) support the hypothesis of a nonlinear relationship between the EUA price and its fundamentals. While they focus on the long-term equilibrium relationship, we turn our attention to the short-term consequences of structural changes in the data-generating process by analyzing data at a daily frequency. This allows for a more profound analysis of short-term fluctuations in expectations and for a more precise estimation. Moreover, we additionally account for potential changes in the volatility of the EUA prices and analyze the entire Kyoto commitment period.[1] Previous literature examining daily data only takes into account selected characteristics of the data or constrains the modeling of structural changes. Alberola, Chevallier, and Chèze (2008) and Chevallier (2009), for example, consider breaks to be deterministic, neglecting the permanently changing nature of the relationship between the EUA price and its fundamentals. Peri and Baldi (2011) consider a fixed threshold that determines the changing impact of crude oil on EUA prices. Benz and Trück (2009), instead, also consider a Markov regime-switching model, but they do not consider the effects of fundamentals on the EUA returns and solely focus on modeling changes in the mean and in the volatility of the EUA price. In this paper, instead, we conduct a combined analysis of the changing nature of the daily price formation process, i.e. we examine the varying relationship between the daily EUA price and its fundamentals and simultaneously allow for changes in volatility. To this end, we estimate Hamilton's (1989) very flexible Markov regime-switching model that is extended by a GARCH structure following Gray (1996) and Klaassen (2002).

In our empirical analysis, we identify two volatility regimes, in which the impacts of the fundamentals differ significantly. Moreover, the probability that the system is in the high volatility regime coincides approximately with the economic recession of 2008 and 2009 and the debt crisis that darkened the economic outlook for Europe in 2011 and 2012. In both periods, economic perspectives and activities were on a decline, leading to a higher uncertainty about the overall stringency of the cap set by European regulators.

The remainder of the paper is organized as follows. Section 5.2 gives a brief overview of

[1] The trial period, i.e. Phase I, is considered as learning period for market participants and the new institutions that frame the EU ETS (Ellerman, Convery, and De Perthuis 2010). The development of trading in Phase I was characterized by periods of low liquidity (Rotfuß 2011). Furthermore, the variation of the price was very low after the price breakdown in 2006. Studies on information processing and market efficiency show that the market was immature during Phase I (e.g. Montagnoli and de Vries, 2010). Since our goal is to analyze the changing nature of the EUA price formation in a mature market environment with well established institutions and experienced market participants, we direct our attention to Phase II. This period is characterized by economic and institutional developments, as for instance the European debt crisis, that enable new insights into the relationship between the EUA price and its fundamentals.

the regulatory design of the EU ETS. Section 5.3 discusses former research on the relation between EUA prices and its fundamentals. Section 5.4 describes the data and Section 5.5 provides the econometric models used in the analysis. Section 5.6 presents the empirical results and Section 5.7 concludes.

5.2 The European carbon market

The EU ETS is one of the key instruments in European climate policy encompassing approximately 50 percent of the total European carbon dioxide emissions. Based on the Directive 2003/87/EC, it was launched in 2005 as the first multinational carbon trading scheme (European Parliament and Council 2003). Designed as cap-and-trade system, it directs pollutant emissions via tradable permits in order to achieve emission reduction targets in a cost-effective and economically efficient way. The regulating institutions set an emission cap for a certain time period - the compliance period - and accordingly allocate a fixed amount of tradable permits among the market participants. Thus, the overall supply of permits is fixed for the considered compliance period. The EU ETS is temporally separated by three compliance periods (Phase I: 2005-2007; Phase II: 2008-2012; Phase III: 2013-2020). Currently, the scheme regulates installations from the power sector and emission-intensive industry sectors such as oil refinement, production and processing of metals, lime, cement, glass, ceramics, chemicals, pulp and paper. In addition to carbon dioxide (CO_2) emissions, the EU ETS covers the greenhouse gases nitrous oxide (N_2O) and perfluorocarbons (PFCs). In 2004, the EU enacted the Linking Directive 2004/101/EC in order to establish a connection between the EU ETS and the project-based mechanisms of the Kyoto protocol (European Parliament and Council 2004): Joint Implementation (JI) and Clean Development Mechanism (CDM). The basic concept is to allow companies to use credits from JI as well as CDM projects, to fulfill their obligations under the EU ETS regulation. The issuance of EUAs takes place gradually, while monitoring, reporting, and verification of the actual emissions as well as the delivery of the equivalent amount of EUAs or credits from project-based mechanisms are executed annually. While the use of EUAs is restricted across compliance periods, banking and borrowing is in principle feasible within each compliance period. This is due to the structure of the yearly iterative time schedule for the allocation and submission of allowances (Ellerman, Convery, and De

Perthuis 2010). The allocation phase for each year ceases in the end of February, while the deadline for the submission of allowances for the previous year is in the end of April, i.e. firms can use permits issued in the current year to comply with the arisen obligations of the previous year.

Phase I is widely seen as the pilot period for newly established institutions and market participants. For Phase I and Phase II, the overall emission cap is defined by the National Allocation Plans (NAPs). The NAPs are determined by each member state and define the national total of permits and the mode of allocation. By approving the NAPs, the European Commission (EC) settles the overall cap. When in April 2006 the information about the actual emissions was released, the market participants began to realize that the overall emission cap for Phase I was not restrictive. Moreover, as neither borrowing nor banking of allowances was allowed between Phase I and Phase II, the price for EUAs issued for Phase I collapsed. The subsequent Phases II and III are connected via banking. Banking of spare allowances extends the time span that is considered by market participants when forming expectations about the overall stringency of the scheme. Thus, banking reduces the exposure and risk of dramatic price drops. Nevertheless, shocks still lead to price adjustments and affect the volatility.

During Phase I and Phase II, the main allocation mechanism was "grandfathering" - the allocation for free, based on historical emissions. In Phase III the EC directly fixes the EU-wide cap without the indirect way of approving NAPs. The allocation mode gradually switches to auctioning as the main allocation mechanism. The cap-setting is stricter (the total amount of 2,04 bn tonnes of carbon dioxide equivalent in 2013 is lowered by 1.74 percent annually until 2020) and more sectors (e.g. production and processing of non-ferrous metals) are regulated since January 2013. For a more detailed description of the changes in Phase III, refer to Directive 2009/29/EC (European Parliament and Council 2009).

Since the introduction of the scheme in 2005, highly efficient EUA spot and derivative markets have evolved. In 2011, the total transaction value in the EU ETS was 122.3 bn euros including credits from the project-based mechanisms (World Bank 2012). The market has been growing rapidly during the first two commitment periods and is now the largest emission market in the world. Several types of transactions and trading products have evolved: EUAs can be traded via bilateral, over-the-counter, or organized markets.

Figure 5.1: EUA price development during Phase II

In addition to the spot market, there is a lively exchange of futures, options, and swaps between the interest groups. Bilateral and over-the-counter transactions dominated trading at exchanges during the first compliance period. Therefore, liquidity at exchanges was low leading to a highly volatile starting of the market (Rotfuß 2011). However, volumes traded on exchanges increased heavily since the beginning of the Phase II. The most liquid derivatives market is situated at the European Climate Exchange (ICE/ECX; London) where approximately 90 percent of the futures contracts are traded. Before its closure in December 2012, the most liquid spot market was Bluenext (Paris). About 70 percent of the daily spot transactions were settled at this exchange.

Figure 5.1 shows the EUA price development during Phase II. After increasing to a peak of about 30 euros in July 2008, the EUA price fell by February 2009 to about 8 euros. This period was characterized by the aftermath of the Lehman Brothers collapse and the subsequent financial turmoils which also caused a slow down of the economic growth in Europe. After a short recovery phase by April 2009, the price followed a lateral movement around 15 euros until June 2011. The following decrease transitioned into a volatile lateral movement during 2012. During this period, again, a weakening economic outlook but also institutional obstacles, such as expected overlapping regulation by the Energy Efficiency Directive (European Commission 2011), hampered the confidence in a restrictive EU ETS cap.

5.3 Related literature

There are several studies that focus on the relation between EUA prices and its determinants. Most of this research is primarily concerned with the existence of various fundamentals and their effects on the EUA price, such as the effects of energy prices, risk factors, or weather conditions. Mansanet-Bataller, Pardo, and Valor (2007), Alberola, Chevallier, and Chèze (2008), and Hintermann (2010) provide evidence for a strong impact of energy prices and extreme temperatures, while Alberola, Chevallier, and Chèze (2009a) and Alberola, Chevallier, and Chèze (2009b) show that also the industrial production of emission intensive sectors affect the EUA price development. Directing the view towards the influence of macroeconomic fluctuations, Chevallier (2009) considers macroeconomic risk factors, which reflect short- and medium-term sentiments in the financial markets about the macroeconomic development. Although macroeconomic risk factors are important determinants for energy commodity futures, their impact on EUA futures appears to be weak. Conrad, Rittler, and Rotfuß (2012) provide evidence that information shocks on regulatory issues and the macroeconomic activity clearly impact EUA prices. According to Anger and Oberndorfer (2008), Oberndorfer and Rennings (2007), Klepper and Peterson (2004), and Demailly and Quirion (2008), the reverse effects of the EU ETS on macroeconomic activity are very weak. The studies of Keppler and Mansanet-Bataller (2010) and Bredin and Muckley (2011) place emphasis on the causal relationships between EUA prices and its fundamentals or their long-term equilibrium relationship. Creti, Jouvet, and Mignon (2012) also contribute to this strand of research. They show, that the equilibrium during Phase II is characterized by an increasing impact of fundamentals on the EUA price in comparison to Phase I.

Overall, the effects of the fundamentals such as energy prices, the weather, the current and future macroeconomic activity, and selected macroeconomic risk factors on carbon prices are clearly evident. The extent and direction of the impact of these fundamentals is, however, not constant over time and highly depends on the sample under consideration. Moreover, structural changes are an important feature of the EUA price generating process. While Alberola, Chevallier, and Chèze (2008) see regulatory announcements as the main reason for those breaks, Chevallier (2009) argues that structural breaks are primarily due to changes in expectations. Recently, Chevallier (2011a), Chevallier (2011b) and Peri and

Baldi (2011) adopt nonlinear models to analyze the long-term equilibrium relationship between the EUA price and its determinants. Their empirical evidence suggests that the European industrial production index and oil prices are likely to influence (eventually asymmetrically and depending on the regime) the EUA price, while reverse effects are not present.

In contrast to Chevallier (2011a) and Chevallier (2011b), we do not rely on monthly data, but exploit daily data of the EUA price and its fundamentals. In doing so, we devote our attention to the short-term consequences of structural changes in the data-generating process. Furthermore, we circumvent the natural loss of information when using aggregated figures or too large intervals between time series observations in order to avoid higher parameter uncertainty, see e.g. Engle (2000). However, daily data exhibits additional features that have to be taken into account, such as conditional heteroscedasticity. Previous studies examining daily data only consider isolated characteristics of the data or impose stark restrictions on structural changes in the relationship between the EUA return series and its fundamentals. Alberola, Chevallier, and Chèze (2008) or Chevallier (2009), for instance, carry out a comprehensive analysis of the relationship between the EUA price and its fundamentals, but they assume breaks to be deterministic and thus do not consider the permanently changing nature of the relationship under investigation. Endogenizing structural changes and allowing for a more flexible approach shall give more robust results. Benz and Trück (2009) adopt such a flexible model by considering a Markov regime-switching model, but they do not include any EUA price determinants into their analysis. We, thus, contribute to the existing literature by simultaneously (i) focusing on the short-term dynamics of the EUA price process, by (ii) modeling changes in the effects of the EUA determinants and in the volatility of the EUA returns via a very flexible Markov regime-switching model that is extended by a GARCH structure, and by (iii) analyzing the entire Kyoto commitment period, i.e. Phase II of the EU ETS.

5.4 Data

Based on the previous studies, we consider several different fundamentals of EUA prices. In the following, we present these fundamentals, the construction of the corresponding series, and provide an analysis of their empirical properties.

5.4.1 EUA prices and their fundamentals

Our empirical analysis exploits data on carbon and energy commodity prices, indicators for macroeconomic risk as well as deviations from the mean temperature in Europe. Our sample period ranges from 1 January 2008 until 31 December 2012, resulting in 1286 daily observations. To obtain a representative carbon price, we follow Chevallier (2009) and use data from the ICE Futures/European Climate Exchange (ECX) which is the most liquid market for carbon derivatives in Europe. We consider annual futures, which expire in December, for the years 2008 up to 2013 and construct the EUA price series based on the daily closing EUA futures prices (EUR/tCO_2e) of the contract with the closest maturity. The prices are used to construct the series of daily continuously compounded returns.[2] The same procedure is applied to energy commodity futures mentioned below. Throughout the paper, we use continuously compounded returns expressed in percentage points.

The link between EUA prices and prices for steam coal, gas, and oil exists mainly because some industries covered by the EU ETS have the ability to switch among various fuels in their production process, see e.g. Mansanet-Bataller, Pardo, and Valor (2007), Alberola, Chevallier, and Chèze (2008), Hintermann (2010) and Creti, Jouvet, and Mignon (2012). Based on different emission and energy intensities, alterations in the price ratio of coal, gas, and oil affect the demand for EUAs and therefore, their price. The fuel switch behavior might cause a reciprocal relationship between carbon and energy commodity prices. However, we assume that the influence of the regionally limited EU ETS on the price formation of the global market for fossil fuels is negligible. In the recent study by Peri and Baldi (2011), this argument finds empirical support. A reliable reference price for European steam coal is the API2 index published by Argus/McCloskey's Coal Price Service. For our investigations, we therefore follow Chevallier (2011b) and employ corresponding API2 index futures prices (USD/t) of annual contracts traded at the European Energy Exchange (EEX). The liquidity of these futures is low due to the fact that a large

[2] This procedure of constructing a price series of a financial asset based on futures contracts with closest maturity is quite common in the literature, see e.g. Bredin and Muckley (2011) and Chevallier (2011b) within the context of EUA prices. It is due to the fact that the contract with the closest maturity date is the most liquid one, such that the construction of the time series is based on the most informative price records.

part of the coal is directly traded via brokers whose transactions are in turn the basis of the API2. Nevertheless, the futures prices are representative, because they are calculated based on fair values enquired from trading members and brokers. Thus, we consider for "Gas" the gas price series (EUR/MWh) based on annual futures contracts traded at the European Energy Derivatives Exchange (ENDEX), which is the largest gas exchange in Europe. Further, we employ the closing prices of the Crude Oil-Brent Current Month Free On Board. Like for the steam coal futures, the price of the crude oil futures is quoted in USD. We use the EUR/USD reference rate published by the European Central Bank (ECB) to convert the coal and oil price series into EUR.

Fama and French (1989) and Sadorsky (2002) have shown the importance of macroeconomic risk factors for the formation of expectations on the equity, bond, and commodity markets. Following Chevallier (2009), we also assume macroeconomic risk factors to influence carbon markets. We expect the EUA price to fall, when the macroeconomic risk measures indicate a prospective economic slow down. This relationship is based on the assumption that adverse business conditions lower aggregated demand and therefore reduce the demand for EUAs. We therefore consider a stock index, a commodity index, and a yield spread as measures for macroeconomic and financial risks. The stock index measures the development of the financial markets and serves as predictor for fluctuations of the overall economic environment. Stock prices reflect expectations about future dividends and can be interpreted as leading indicators for the development of business conditions.[3] We include into our analysis the Dow Jones EURO STOXX 50 (DJES50) which represents 50 European companies that are leading in their industries.[4] The index includes different branches such as energy generation, manufacturing, and financial institutions. We further consider an indicator capturing risk related to fluctuations at the global commodity markets, i.e. the Thomson Reuters/Jeffries Commodity Research Bureau Index (CRBI), which

[3] Fama and French (1987), Sadorsky (2002) and Chevallier (2009) exploit dividend yields as leading indicator for economic activities. Following Bredin and Muckley (2011), we consider the stock index itself instead of including dividend yields. Based on the dividend discount model, the stock index itself can also be interpreted as a leading indicator.

[4] Note that the empirical literature does not provide clear evidence on the lead-lag relationship of futures versus spot prices for the DJES50 index. Robles-Fernandez, Nieto, and Fernandez (2004), for example, show that the information flow between the DJES50 futures and spot prices is bidirectional. We have therefore decided to use the spot prices as this is a highly liquid market.

reflects the price development of several commodity classes. The prices of commodities are expected to decrease in times of lowering economic activity induced by decreasing aggregated demand (Chevallier 2009).

To account for default risks in credit markets, we include the default spread defined as the difference between two yields to maturity of two fixed income portfolios which represents the premium required to compensate a lender for investing in the riskier asset. In our case, we use data of average annual yields of U.S. corporate long-term bonds rated AAA and BAA, that are published by Moody's. Empirical findings by Fama and French (1989) provide evidence that the default spread rises in times of high economic uncertainty. The findings are congruent with the development of the default spread in our data set. Thus, we expect the EUA price to decline, when the default spread increases.

Similar to Mansanet-Bataller, Pardo, and Valor (2007) and Hintermann (2010), we also include variables reflecting extreme weather conditions into our analysis. In particular, we consider deviations from average temperatures. The deviation is computed as difference between the daily measured temperature and the mean of the monthly temperature averages over the years from 2005 up to 2012. The basis for the series, to which we simply refer as "Temperature", is the Tendances Carbone European Temperature Index, which is obtained as the weighted average of the daily European temperatures.[5] The weights are the shares of the NAPs in the considered countries. Extremely high or low temperatures increase the demand for heating or cooling and raise therefore emissions as well as EUA prices. For our empirical analysis, we consider values of the deviations from average temperatures.

5.4.2 Empirical analysis

Table 5.1 highlights the empirical properties of the employed data by presenting the descriptive statistics of the continuously compounded return series and of the deviations from average temperature. Obviously, the mean and median values of all return series are

[5]The Tendances Carbone European Temperature Index is computed as NAP weighted average of the daily temperature of the Metnext Weather Indices of 4 countries (German, Italy, France, and UK) from January 2005 to September 2009 and of temperature data of 18 countries since October 2009. The Metnext Weather Indices are intra-country temperature averages weighted by population. We are grateful to CDC Climat research, Climpact Metnext for kindly providing us with this data.

Table 5.1: Descriptive statistics

Variable	Mean	Median	St. Dev.	Skewness	Kurtosis	Min	Max	Jarque-Bera
EUA	-0.1132	0.0000	2.6400	-0.2414	4.8933	-11.6029	11.3659	204
Oil	0.0155	0.0662	2.0587	-0.0473	5.6069	-8.4875	9.3543	364
Coal	-0.0136	-0.0194	1.5748	-0.3514	8.7243	-10.1216	8.5801	1782
Gas	-0.0291	-0.0255	1.4831	0.3999	6.4272	-7.4032	9.2129	663
DJES50	-0.0407	-0.0352	1.8027	0.1061	7.3007	-8.2078	10.4376	993
CRBI	0.0103	0.0356	0.5977	-0.7415	7.2826	-3.5645	2.3149	1101
Default Spread	-0.0123	0.0000	1.8034	-0.0557	13.2895	-13.3532	14.1078	5674
Temperature	0.0120	0.1049	2.5507	-0.3755	3.2681	-9.7159	7.1599	34

Notes: Reported are the descriptive statistics of the daily return series and of the daily deviations from average temperature. All returns are continuously compounded and expressed in percentage points. For further information on the variables and their construction, we refer to Section 5.4.1. We also report the individual Jarque-Bera test statistics on normality of each series (last column). The corresponding critical value at the 5 percent significance level is 5.99.

very small. The same is true for the temperature deviations. The mean is in all cases not significantly different from zero. All time series clearly exhibit excess kurtosis. The Jarque-Bera test rejects the null hypothesis of zero skewness and excess kurtosis for all return series as well as for the deviations from average temperature. According to the test statistics of the Augmented Dickey-Fuller (ADF) test and the Phillips-Perron test (PP), which are reported in Table 5.2, the logarithms of energy prices and the prices of the macroeconomic and financial risk factors are nonstationary, while the test results indicate that the corresponding continuously compounded returns are stationary.[6] Applying the tests to the deviations from average temperatures shows, that these are stationary and we therefore do not compute unit root tests on the first differences of this series.

In the remainder of the paper, we consider for our empirical analysis the stationary time series, i.e. the continuously compounded returns of futures prices and risk factors as well as deviations from average temperatures. Figure 5.2 depicts the time evolvement of the employed series. All return series show the well-known phenomena of volatility

[6] Using prices constructed from rolling over between futures contracts may induce jumps, which in turn may affect inference based on unit root tests. We do not explicitly account for these jumps when conducting the unit root tests on the price series, but we expect that the jumps do not affect our results. The study of Medina and Pardo (2012) supports this. In particular, analyzing separately the price series of individual futures contracts, they find that the EUA prices are integrated of order one, which is consistent with our unit-root test results.

Figure 5.2: Return series

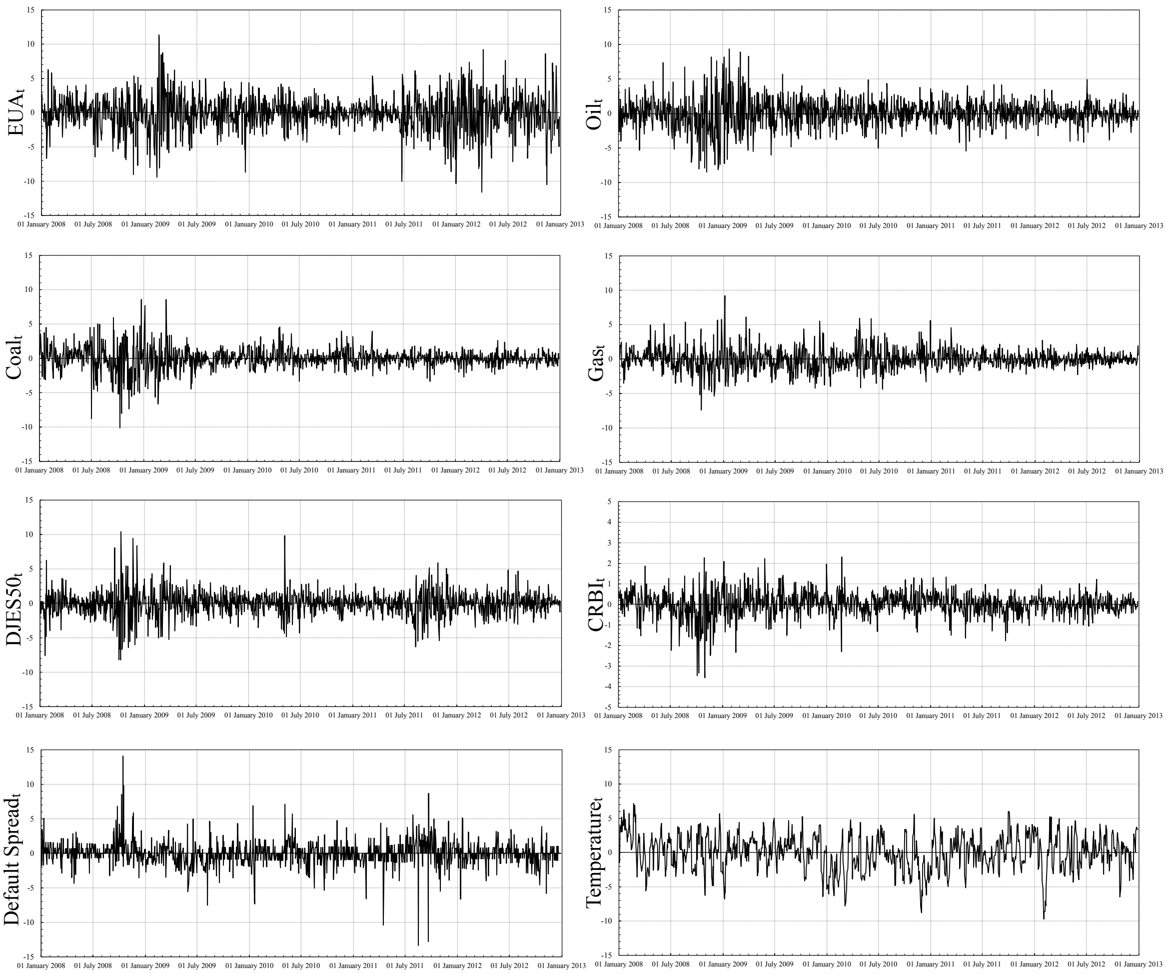

Notes: Time evolvement of the daily returns, and of the daily deviations from average temperature (bottom right panel). All returns are continuously compounded and expressed in percentage points. For further information on the variables and their construction, we refer to Section 5.4.1.

Table 5.2: Unit root tests

	Augmented Dickey-Fuller test				Phillips-Perron test			
	p_t		r_t		p_t		r_t	
Variable	test statistic	p-value	test statistic	p-value	test statistic	p-value	test statistic	p-value
EUA	-1.550	0.5089	-33.382	0.0000	-1.462	0.5521	-33.318	0.0000
Oil	-1.102	0.7144	-36.585	0.0000	-1.088	0.7199	-36.578	0.0000
Coal	-1.724	0.4188	-24.070	0.0000	-1.818	0.3713	-32.996	0.0000
Gas	-1.343	0.6095	-23.710	0.0000	-1.353	0.6046	-31.403	0.0000
DJES50	-3.191	0.0862	-36.797	0.0000	-3.128	0.0997	-37.043	0.0000
CRBI	-0.977	0.7614	-12.199	0.0000	-0.859	0.8011	-32.628	0.0000
Default Spread	-1.501	0.5332	-11.120	0.0000	-1.043	0.7372	-36.471	0.0000
Temperature	-11.927	0.0000	-	-	-11.662	0.0000	-	-

Notes: Reported are the test statistics and the corresponding p-values of the Augmented Dickey-Fuller test and the Phillips-Perron test on the null of a unit root in the logarithmic price series (columns 2-3 and 6-7), in the corresponding continuously compounded returns (columns 4-5 and 8-9), as well as in the daily deviations from average temperature (last row). Note, as the null of a unit root is rejected for the deviations from average temperature, we do not report unit root tests for the first differences of this series. The variables are defined in Section 5.4.1.

clusters. Moreover, during the financial crisis or more precisely during the aftermath of the Lehman Brothers bankruptcy in September 2008, all return series exhibit high levels of volatility. The sample autocorrelation function of the squared EUA returns is depicted in Figure 5.3 along with the corresponding 95 percent Bartlett confidence intervals. The sample autocorrelation function of the squared returns is slowly decaying, which is typical for daily returns exhibiting volatility clustering.[7] This pattern has also already been observed for emission allowance returns by Paolella and Taschini (2008), and Medina and Pardo (2012). The presented empirical properties of EUA returns are thus in line with the stylized facts of financial asset returns. This is also reported in Medina and Pardo (2012), who conduct a more detailed analysis of the properties of the EUA returns and find that the EUA returns additionally exhibit features that are common to commodity futures.[8]

[7] In the levels, the EUA returns exhibit significant autocorrelation only at the first lag. The first order autocorrelation is 0.0706.

[8] An extensive empirical analysis of our data with respect to these features is out of the scope of the present paper.

Figure 5.3: Sample autocorrelation function of the EUA returns

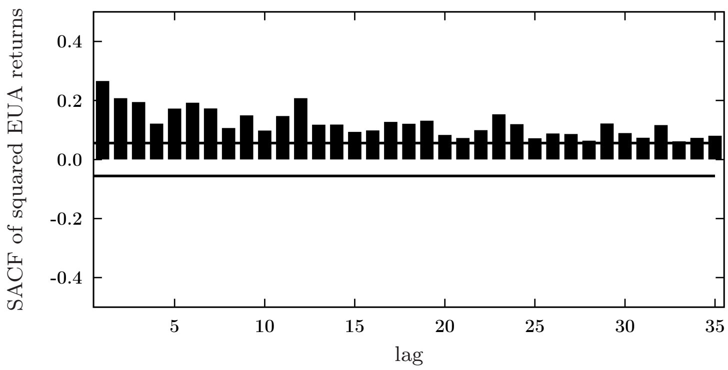

Notes: Sample autocorrelation function of the daily squared continuously compounded EUA returns (expressed in percentage points).

5.5 Methodology

We model the changing nature of the relation between the EUA return and its fundamentals via a Markov regime-switching model (Hamilton 1989). The model is very flexible and, thus, able to describe quite complex dynamics of the considered time series. The model used in this paper is based on Gray (1996) and Klaassen (2002), who extend Hamilton's original approach by including a generalized autoregressive conditional heteroscedasticity (GARCH) structure into the Markov regime-switching model. The model thus allows for structural changes in financial volatility, which may generate the observed persistence in volatility, see e.g. Diebold and Inoue (2001), Granger and Hyung (2004) and Mikosch and Stărică (2004). Furthermore, Paolella and Taschini (2008), Chevallier (2009), and Medina and Pardo (2012) find, that EUA return series based on daily data show characteristics such as volatility clustering and fat tails. They therefore advocate the use of GARCH-type models in order to take these stylized facts of asset returns into account when considering EUA return series.

The model is based on the assumption that the data generating process shifts at different points of time and that these discrete aperiodic shifts between a finite number of *states* or *regimes* are driven by a hidden Markov chain. In the following, we briefly explain the model structure in more detail. To this end, let r_t denote the daily continuously compounded EUA return at time period t ($t = 1, 2, \ldots, T$) and let $s_t \in \{0, 1\}$ be a latent *state*-variable that governs the switch between two possible regimes. The EUA returns are assumed to be affected by k fundamentals, which are subsumed in the vector $\boldsymbol{x}_t' = (1, x_{1t}, x_{2t}, \ldots, x_{kt})$. The influence of these fundamentals on the EUA returns is allowed to vary over time, which is highlighted by the superscript s_t on the parameter vector $\boldsymbol{\gamma}^{(s_t)}$. In particular, we assume that the impact depends on the current state s_t. The mean equation of our model is therefore given by

$$r_t = \boldsymbol{x}_t' \boldsymbol{\gamma}^{(s_t)} + a_t. \tag{5.1}$$

where the parameter vector $\boldsymbol{\gamma}^{(s_t)} = (\gamma_0^{(s_t)}, \gamma_1^{(s_t)}, \gamma_2^{(s_t)}, \ldots, \gamma_k^{(s_t)})'$ with $s_t \in \{0, 1\}$ measures the influence of the risk factors on EUA returns in the two regimes.

To account for the possibility of structural changes in the volatility process, we follow Klaassen (2002) and assume that the innovation a_t is normally distributed with zero mean and variance $Var(a_t|\tilde{s}_t, I_{t-1})$ conditional on the regime path $\tilde{s}_t = (s_t, s_{t-1}...)$ and the information set I_{t-1} containing the information available at time $t-1$:

$$a_t|\tilde{s}_t, I_{t-1} \sim N(0, Var(a_t|\tilde{s}_t, I_{t-1})). \tag{5.2}$$

The dynamics of the conditional variance $Var(a_t|\tilde{s}_t, I_{t-1})$ is based on a GARCH(1,1) model, where, however, the parameters of the conditional variance equation are also allowed to be state dependent, i.e. the conditional variance is given by

$$Var(a_t|\tilde{s}_t, I_{t-1}) = \omega^{(s_t)} + \alpha^{(s_t)} a_{t-1}^2 + \beta^{(s_t)} E[Var(a_{t-1}|\tilde{s}_{t-1}, I_{t-2})|\tilde{s}_t, I_{t-1}]. \tag{5.3}$$

That is, in each regime the conditional variance is given by a regime-specific GARCH(1,1) model. Similarly to the single-regime GARCH(1,1) model, the regime-specific unconditional variance is then given by $\sigma^{2(s_t)} = \left(1 - \alpha^{(s_t)} - \beta^{(s_t)}\right)^{-1} \omega^{(s_t)}$ for $s_t = 0, 1$.[9]

The development of s_t and therefore, the switching in regimes is governed by a homogenous first order Markov chain and can be fully described by the *transition probabilities* p and q which refer to the probabilities of being in the same state s_t as in the previous period, i.e.

$$P[s_t = 1|s_{t-1} = 1] = p, \qquad P[s_t = 0|s_{t-1} = 0] = q. \tag{5.4}$$

[9]Note that in our empirical analysis we have also experimented with alternative lag orders of the GARCH(p, q) model, but the GARCH(1,1) model is preferred by both, the Akaike and the Bayesian information criteria. So we consider here the GARCH(1,1) model, i.e. a lag order specification that is also standard in the GARCH literature, see e.g. Medina and Pardo (2012) and Paolella and Taschini (2008) within the context of EUA returns. The results of this preliminary analysis are available upon request.

Within each regime the relationship between the EUA return and its fundamentals is linear (see Equation 5.1) and the state variable s_t, thus, governs the shift between these two linear relationships. The transition probabilities characterize the switching in regimes and therefore the evolvement of the system over time. In order to draw inference on s_t, we calculate the smoothed probabilities $P[s_t = j|I_T]$, $j = 0, 1$ based on the algorithm provided by Kim (1994). The smoothed probabilities are conditional on the information set I_T that comprises the entire information contained in the data set.

5.6 Empirical results

5.6.1 Non-switching GARCH model

In accordance to former research on the relation between the EUA return and its fundamentals, we begin our empirical analysis by estimating a GARCH(1,1) model without regime-switching. This allows us to compare our empirical findings to the existing empirical literature and to highlight special features. To this end, we regress the EUA returns on selected energy variables, a stock index, a commodity index, a default spread and a temperature variable reflecting extreme weather conditions in Europe. The model, to which we simply refer to as the "non-switching GARCH model", takes the following form

$$\begin{aligned} \text{EUA}_t &= \gamma_0 + \gamma_1 \text{Oil}_t + \gamma_2 \text{Coal}_t + \gamma_3 \text{Gas}_t \qquad (5.5) \\ &\quad + \gamma_4 \text{DJES50}_t + \gamma_5 \text{CRBI}_t + \gamma_6 \text{Default Spread}_t \\ &\quad + \gamma_7 \text{Temperature}_t + a_t, \quad \text{where} \\ a_t|I_{t-1} &\sim N(0, Var(a_t|I_{t-1})) \quad \text{and} \\ Var(a_t|I_{t-1}) &= \omega + \alpha a_{t-1}^2 + \beta Var(a_{t-1}|I_{t-2}), \end{aligned}$$

i.e. we account for conditional heteroscedasticity by specifying a GARCH(1,1) model.[10] We estimate this model for the data analyzed in Section 5.4.2. In particular, we use

[10]Note that we have estimated a linear regression model without allowing for GARCH effects as a first step. However, the Lagrange multiplier test rejects the null hypothesis of no ARCH effects in the innovations of this specification. This supports once more the necessity to account for volatility clustering when modeling daily EUA returns, see also our results in the empirical data analysis of Section 5.4.2. The results of this preliminary analysis are available upon request.

Table 5.3: Estimation results of non-switching GARCH(1,1) model

Variable	Parameter	Std. error
Constant	−0.0235	(0.0548)
Oil	0.0946***	(0.0318)
Coal	0.0739**	(0.0373)
Gas	0.3536***	(0.0374)
DJES50	0.2202***	(0.0300)
CRBI	0.2115**	(0.1051)
Default Spread	0.0068	(0.0300)
Temperature	0.0183	(0.0205)
ω	0.1553***	(0.0369)
α	0.1437***	(0.0164)
β	0.8365***	(0.0169)

Notes: The table presents the estimation results of the GARCH(1,1) model without regime-switching as given in Equation (5.5). *** indicates significance at 1 percent, ** at 5 percent, * at 10 percent.

the daily continuously compounded returns (expressed in percentage points) of the EUA price and of its fundamentals, with the exception of the extreme weather variable, which is measured as daily deviations from average temperature. The specific variables are explained in detail in Section 5.4.1.

The estimation results of the non-switching GARCH(1,1) model are reported in Table 5.3. With respect to the energy variables, our results are primarily in line with the existing literature summarized in Section 5.3. More specifically, just like in the existing literature, we find a significant and positive impact of crude oil and natural gas on the EUA returns. We further observe a positive and significant effect of coal, which has recently also been documented in Chevallier (2011b) using data on Phase I and part of Phase II. In contrast, Alberola, Chevallier, and Chèze (2008) and Aatola, Ollikainen, and Toppinen (2013) report a significant, but negative influence on the EUA returns based on samples starting in January 2005 and lasting until December 2008, and December 2010, respectively. Focusing exclusively on Phase I, Hintermann (2010) finds no significant effect of coal. Thus, the empirical evidence on the effect of coal on EUA returns is mixed, which may be due to two, potentially offsetting, effects that determine the relationship between coal and EUA prices. First, it is widely accepted that aggregated economic activity can drive the demand for commodities and raw materials and, thus, their prices. Increasing or decreasing aggregated demand might effect commodity and EUA prices in the same way, leading to a positive relationship also in the returns. Second, some sectors might have the possibility to substitute coal by combusting gas or other fuels. This fuel-switch behavior implies that in a situation of increasing coal prices, a company will ceteris paribus switch to less expensive fuels and in the case of oil and gas, less emission-intensive fuels, leading to a negative relationship between coal and EUA returns. Our finding of a significant, positive effect of coal suggests that the aggregated demand effect overweighs the substitution effect during Phase II.

Our estimation results also suggest that the EUA returns are affected by macroeconomic and financial risk factors. In particluar, the stock index and the commodity index have a positive and significant effect at the 5 percent significance level. The positive impacts of the stock index and the commodity index are consistent with our expectations: Market participants associate a positive development of stock index values or commodity prices with rising economic activity, which leads to increasing EUA prices. The positive impact of the stock index is also in line with the empirical results of Chevallier (2009) and Medina and Pardo (2012).[11] Using a different dataset, Hintermann (2010) finds instead no significant impact. The insignificance of the default spread is also documented in Chevallier (2009).

The coefficient of the temperature variable is not statistically different from zero. The parameter estimates of the volatility equation are consistent with those that are commonly observed when fitting a GARCH(1,1) model to daily financial return series indicating a highly persistent volatility of the EUA returns.

The comparison of the results of the existing literature suggests, that the relation between the EUA price and its fundamentals is time-varying. Even within comparable model classes, the impacts heavily depend on the time periods considered. We therefore explicitly take into account the changing nature of this relationship by estimating a Markov regime-switching model.

[11] Moreover, Chevallier (2009) reports a negative impact of dividend yields on EUA prices, which is also in accordance with our positive impact of the stock index, as dividend yields are reciprocal to the values of the corresponding stock index.

5.6.2 Markov regime-switching GARCH model

We now turn our attention to the empirical analysis of the Markov regime-switching model discussed in Section 5.5, in which, both, the effects of the fundamental factors on EUA returns and the volatility dynamics are allowed to depend on two regimes. Preliminary estimation results of this model (not reported here), however, show that a GARCH(1,1) specification is not needed in the second regime. In particular, the GARCH parameters $\alpha^{(1)}$ and $\beta^{(1)}$ are insignificant, implying that the second regime is characterized by a constant volatility. We therefore exclude the GARCH(1,1) specification from the second regime by restricting $\alpha^{(1)}$ and $\beta^{(1)}$ to zero, such that, in the second regime, the unconditional variance of the shocks to the EUA returns is given by $\sigma^{2(1)} = \omega^{(1)}$. For the first regime, instead, we keep the GARCH(1,1) specification, as $\alpha^{(0)}$ and $\beta^{(0)}$ are both significant at the 5 percent significance level. The unconditional variance of the first regime is therefore still given by $\sigma^{2(0)} = \left(1 - \alpha^{(0)} - \beta^{(0)}\right)^{-1} \omega^{(0)}$.[12] The estimation results of this restricted Markov regime-switching model are presented in Table 5.4. The smoothed probabilities of being in Regime 1 are depicted in Figure 5.4.

Both regimes are characterized by no clear price trend. However, the first regime depicts a high unconditional variance of the shocks to the EUA returns, $\sigma^{2(0)}$, and can be related to phases of higher uncertainty. In both regimes, the gas and the equity index returns individually have a significant impact on EUA returns, while the parameters associated with extreme temperatures and with the commodity index are insignificant. Overall, we do not observe huge differences in the estimates of the parameters in the mean equations of the two regimes. Thus, in order to rule out, that regime switches are only relevant for the volatility dynamics of EUA returns, we conduct a Likelihood-ratio test on the null hypothesis, that the parameters of the mean equation are identical across the two regimes, i.e. we test H_0: $\boldsymbol{\gamma}^{(0)} = \boldsymbol{\gamma}^{(1)}$ against the alternative H_1: $\boldsymbol{\gamma}^{(0)} \neq \boldsymbol{\gamma}^{(1)}$. According to the test statistic, which is also reported in Table 5.4, the null hypothesis is rejected at the 5 percent significance level. Thus, we conclude that the impact of fundamentals on the EUA price is state dependent. In the first regime, the most important drivers of the EUA returns are the returns of the equity index and of gas. The returns of coal have a weaker, but significant, impact which is insignificant in the second regime (at the 5 percent

[12]The results of this preliminary analysis are available upon request.

Table 5.4: Estimation results of the Markov regime-switching GARCH model

Variable	Regime 1 ($s_t = 0$) Parameter	Std. error	Regime 2 ($s_t = 1$) Parameter	Std. error
Constant	−0.1285	(0.0934)	0.0737	(0.0632)
Oil	0.0846*	(0.0509)	0.1205***	(0.0428)
Coal	0.1645**	(0.0693)	−0.0955*	(0.0514)
Gas	0.3098***	(0.0748)	0.3794***	(0.0440)
DJES50	0.2657***	(0.0559)	0.1538***	(0.0446)
CRBI	0.2999	(0.1869)	0.1293	(0.1285)
Default Spread	−0.0380	(0.0511)	0.0473	(0.0345)
Temperature	0.0243	(0.0388)	0.0122	(0.0273)
ω	1.0923***	(0.3221)	-	-
α	0.1463***	(0.0361)	-	-
β	0.7251***	(0.0645)	-	-
Uncond. variance of EUA return shocks	$\sigma^{2(0)}$=8.4920		$\sigma^{2(1)}$=1.1516	
Transition probabilities	P[$s_t = 0 \vert s_{t-1} = 0$]=0.9837		P[$s_t = 1 \vert s_{t-1} = 1$]=0.9676	
LR test[a]	Test statistic		Crit. value	
H_0: $\boldsymbol{\gamma}^{(0)} = \boldsymbol{\gamma}^{(1)}$ H_1: $\boldsymbol{\gamma}^{(0)} \neq \boldsymbol{\gamma}^{(1)}$	15.821		15.507	

Notes: The table shows the estimation results of the Markov regime-switching GARCH model. *** indicates significance at 1 percent, ** at 5 percent, * at 10 percent. The lower panels of the table report the estimates of the transition probabilities, the unconditional variance of the shocks to the EUA returns in both regimes, and the result of a Likelihood Ratio test for the H_0 hypothesis, that the parameters of the mean equation are identical across the regimes.

[a] The likelihood ratio test statistic is χ^2 distributed with 8 degrees of freedom. The reported critical value corresponds to the 5 percent significance level.

significance level). In the second regime, the oil and the equity index returns each have modest impacts on the EUA return, while the strongest impact is again observed for gas.

The significant effects of the energy variables (coal, gas, and oil) are positive. Moreover, the equity index has a significant and positive impact on EUA returns in both regimes. These findings are in line with our results based on the non-switching GARCH model and with the economic intuition and the literature discussed in Section 5.6.1. Moreover, the strong impact of the equity index reflects its importance as a predictor for the general economic development and, thus, for the aggregated demand for allowances. In accordance with the results of Chevallier (2009), we find no significant impact of the commodity index nor of the interest rate spread on the EUA returns. The coefficient of the temperature variable, which reflects extreme weather conditions, is also statistically insignificant. The existing literature, that studies the impact of extreme temperatures and weather events on EUA returns, provides mixed evidence and the results heavily depend on the considered sample periods. E.g., while Hintermann (2010) does not find a significant impact of extreme temperatures either, Alberola, Chevallier, and Chèze (2008) report that extreme temperatures had a significant impact during the strong winter seasons in 2006 and 2007 depending on the underlying subsample analysis.

Furthermore, according to the estimation results reported in Table 5.4, the conditional volatility in the first regime is persistent, as measured by $\alpha^{(0)} + \beta^{(0)} = 0.8714$, but smaller than the estimated persistence for the non-switching GARCH model, which is 0.9802. This is consistent with the observation that persistence in volatility may be generated by less persistent processes with structural breaks or regime switches, see e.g. Chen, Härdle, and Pigorsch (2010), Diebold and Inoue (2001) and Granger and Hyung (2004).[13] Moreover, the transition probabilities from the first to the second regime and vice versa are very small. In other words, switching from one regime to the other is not very likely and

[13] Recall that based on our preliminary estimation results, we model the volatility in the second regime as a constant. As such, there is no specification for the persistence of the volatility in this regime. Comparing the time evolvement of the EUA returns with the smoothed probabilities, both depicted in Figure 5.4, further indicates, that the volatility in the second regime indeed seems to stay at the same level and that there is no indication of volatility clustering. This is also observed for the smoothed probabilities of the Markov regime-switching-GARCH model without restricting the volatility of the second regime to be constant (the figure is not presented here, but available upon request), which is in line with the observation of the insignificant parameter estimates $\alpha^{(1)}$ and $\beta^{(1)}$.

changes in the relation between the EUA returns and the considered fundamentals do not occur very frequently.

Figure 5.4 depicts the EUA price and its continuously compounded return series along with the time evolvement of the conditional variance of the first regime and the smoothed probabilities of being in the first regime. The figure illustrates, that after a short consolidation phase in January 2008, the carbon price series shows an upward trend during the first half of the year 2008. Except for the consolidation phase, the smoothed probabilities presented in the bottom panel of Figure 5.4 indicate the system to be with high probability in the second regime that is characterized by a low and constant volatility. This is also in line with the behavior of the carbon return series during this year. The smoothed probabilities suggest that the occurrence of the first regime coincides approximately with the economic recession of 2008 and 2009. In this period, economic activity slowed down and lowered the demand for commodities and EUAs. As a consequence, commodity and allowance prices decreased dramatically. The return series r_t and the conditional volatility $Var(a_t|s_t = 0, I_{t-1})$ in Figure 5.4 show that the uncertainty in the market increased in the second half of 2008 and the first half of 2009. In the time period from early 2009 to August 2010, the carbon price is at a level of around 15 euros. During this phase of lateral movement, the smoothed probabilities still indicate that the first regime prevails. The conditional volatility is still on higher levels in comparison to the phase of the upward trend in the first half year of 2008. Also, the second half year of 2010 is characterized by a lateral movement ending up in a slight upward movement during the first half year of 2011. The smoothed probabilities indicate the model to be in the second regime during this period, this coincides with a phase of economic recovery. Our results indicate, that the recovery after the economic cool down stabilized the expectations about the stringency of the EU ETS leading to slightly increasing prices and a decreasing level of uncertainty. In June 2011, there is a sharp drop in EUA prices, followed by a downward movement until early 2012. The tremendous fall in June occurred directly after the announcement of the Energy Efficiency Directive (European Commission 2011). The proposed directive includes energy efficiency goals, which imply decreasing emissions that were not taken into account by the cap setting of the EU ETS. From this date onwards, the smoothed probabilities suggest the first regime to be dominant until the end of 2012 accompanied with high levels of volatility. There are several reasons that caused this downward movement

Figure 5.4: Inference about the latent state variable: smoothed probabilities

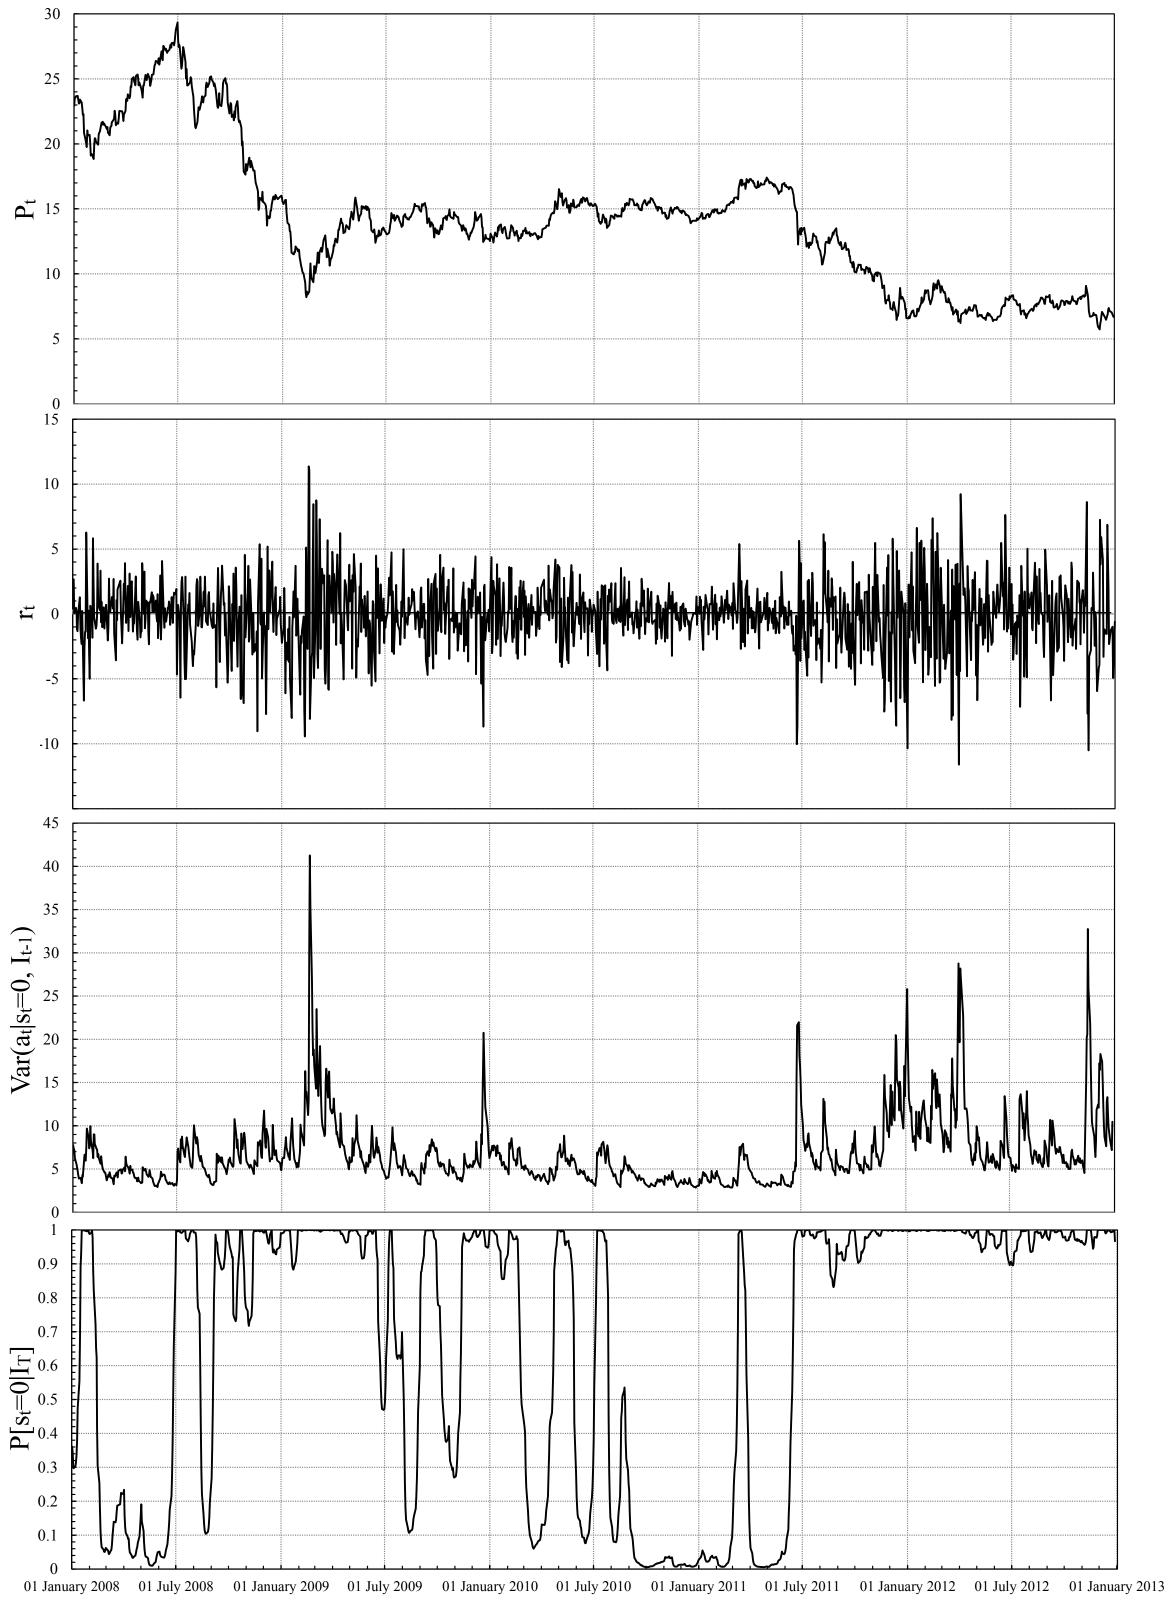

Notes: Inference about the regimes s_t. Depicted are the time evolvement of the EUA prices, P_t, the corresponding returns (continuously compounded and expressed in percentage points), r_t, the conditional variance of the first regime $Var(a_t|s_t = 0, I_{t-1})$, and the smoothed probabilities for being in the first regime $P[s_t = 0|I_T]$. Further details on the construction of the price and return series as well as on the definitions of the conditional variance and the smoothed probabilities are provided in Sections 5.4.1 and 5.5, respectively.

and the high uncertainty in the market. First, the full extent of the economic slowdown due to the financial crisis and the subsequent European debt crisis became clearer after the data on historic verified emissions discharged in 2010 were published in 2011. The data show that the short economic recovery after the slowdown of the financial crisis did not reduce the oversupply of EUAs significantly. In contrast, the emissions data provided evidence that huge amounts of EUAs are banked to Phase III also due to the subsequent European debt crisis and the associated decrease in economic activities in Europe (World Bank 2012). Second, the expectation among market participants arose, that the weakening growth outlook for the EU in the early Phase III might hamper the stringency of the system even further.[14] Third, on the regulatory side, overlapping regulation, the lack of adjustment of the 20 percent reduction goal for 2020, and the weakness in international ambition for reducing greenhouse gases supported the expectation of the market participants that the cap set for Phase II and Phase III will not be restrictive. During 2012, the price moves laterally around a level of 7 euros, while the volatility stays on high levels.

Our findings suggest, that the strong economic slowdown during the aftermath of the financial crisis, the weakening growth outlook due to the European debt crisis and regulatory obstacles coincide with the state at highly volatile EUA prices. In this kind of situation, the market participants are uncertain about the overall supply and demand ratio. They fear the market condition of an oversupply of allowances, where the overall cap for the second and third trading period is not binding anymore resulting in a sharp price drop. This uncertainty seems to increase in times, when the expectations about the general economic development deteriorates or information is released that indicates less stringency for the EU ETS.

5.7 Conclusions

This paper is concerned with the nonlinear relationship between the EUA price and its fundamentals. We argue that changes in the data generating process of the EUA price are a consequence of the design of the EU ETS. In particular, since the EU ETS runs on the basis of a cap-and-trade system, the supply of allowances is fixed over a certain period of time,

[14] Indicators of the current and expected economic development, e.g. the CESifo World Economic Survey, clearly display the weakening economic outlook back in July 2011 (Plenk, Nerb, Wohlrabe, and Berg 2013).

while the demand may be affected by both positive and negative economic shocks. When negative shocks reduce current emissions and, thus, the current demand for allowances, uncertainty about the overall stringency of the scheme increases and market participants adjust their expectations. The associated EUA trading translates into higher volatility and a varying relation between the EUA price and its fundamentals. Our empirical results support such a nonlinearity in the dynamics of the EUA price. We estimate a Markov regime-switching GARCH model, accounting for changing states in the mean and variance of the EUA returns. Our model is able to identify a low and a high volatility regime and shows significant differences in the impact of the fundamentals across states. The high volatility regime largely coincides with phases when weakening economic conditions or institutional changes impair the confidence in the stringency of the cap set in accordance with the EU emission targets. In 2008 and 2009, when the overall actual emissions were on a decline due to the economic recession caused by the financial crisis, our model indicates the high volatility regime to be predominant. This also applies for the time period from July 2011 until December 2012, when the debt crisis weakened the economic outlook for Europe and institutional announcements hampered the confidence in a stringent EU ETS cap.

5.8 Acknowledgements

We are grateful to the editor and three anonymous referees for their valuable comments. Uta Pigorsch thanks the Ministerium für Wissenschaft, Forschung und Kunst des Landes Baden Württemberg for financial support through the Juniorprofessorenprogramm im Rahmen der Zukunftsoffensive IV. Moreover, we thank CDC Climat research, Climpact Metnext for kindly providing us with the Tendances Carbone European Temperature Index data. The usual disclaimer applies.

Bibliography

Aatola, P., M. Ollikainen, and A. Toppinen (2013): "Price determination in the EU ETS market: Theory and econometric analysis with market fundamentals," *Energy Economics*, 36, 380–395.

Abadie, A., and G. W. Imbens (2006): "Large Sample Properties of Matching Estimators for Average Treatment Effects," *Econometrica*, 74, 235–267.

Abrell, J., A. N. Faye, and G. Zachmann (2011): "Assessing the impact of the EU ETS using firm level data," Bruegel Working Paper.

Ackerberg, D., C. L. Benkard, S. Berry, and A. Pakes (2007): "Econometric Tools for Analyzing Market Outcomes," *Handbook of econometrics*, 6, 4171–4276.

Ackerberg, D., K. Caves, and G. Frazer (2015): "Identification Properties of Recent Production Function Estimators," *Econometrica*, 83, 2411–2451.

AG Energiebilanzen e.V. (2014): "Energiebilanz der Bundesrepublik Deutschland 2003 - 2012," http://www.ag-energiebilanzen.de/7-0-Bilanzen-1990-2014.html, extracted on 29th of April, 2016.

Aigner, D., C. K. Lovell, and P. Schmidt (1977): "Formulation and estimation of stochastic frontier production function models," *Journal of Econometrics*, 6, 21–37.

Alberola, E., J. Chevallier, and B. Chèze (2008): "Price Drivers and Structural Breaks in European Carbon Prices 2005 - 2007," *Energy Policy*, 36, 787–797.

——— (2009a): "The EU Emissions Trading Scheme: The Effects of Industrial Production and CO2 Emissions on European Carbon Prices," *International Economics*, 116, 95–128.

——— (2009b): "Emissions Compliances and Carbon Prices under the EU ETS: A Country Specific Analysis of Industrial Sectors," *Journal of Policy Modeling*, 31, 446–462.

Aldy, J. E., A. J. Krupnick, R. G. Newell, I. W. Parry, and W. A. Pizer (2010): "Designing climate mitigation policy," *Journal of Economic Literature*, 48, 903–934.

Anger, N., and U. Oberndorfer (2008): "Firm Performance and Employment in the EU Emissions Trading Scheme: An Empirical Assessment for Germany," *Energy Policy*, 36, 12–22.

Angrist, J. D., and J.-S. Pischke (2009): *Mostly harmless econometrics: an empiricist's companion*, vol. 1. Princeton University Press, Princeton.

Angrist, J. D., and J.-S. Pischke (2010): "The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics," *Journal of Economic Perspectives*, 24, 3–30.

Arlinghaus, J. (2015): "Impacts of Carbon Prices on Indicators of Competitiveness: A Review of Empirical Findings," OECD Environment Working Paper No. 87, OECD Publishing ENV/WKP(2015)87, Paris.

Aw, B. Y., M. J. Roberts, and D. Y. Xu (2011): "R&D Investment, Exporting, and Productivity Dynamics," *American Economic Review*, 101, 1312–1344.

Balcilar, M., R. Demirer, S. Hammoudeh, and D. K. Nguyen (2016): "Risk spillovers across the energy and carbon markets and hedging strategies for carbon risk," *Energy Economics*, 54, 159–172.

Becker, R., and V. Henderson (2000): "Effects of air quality regulations on polluting industries," *Journal of political Economy*, 108, 379–421.

Benz, E., and S. Trück (2009): "Modeling the Price Dynamics of CO2 Emission Allowances," *Energy Economics*, 31, 4–15.

Berman, E., and L. T. Bui (2001): "Environmental regulation and productivity: evidence from oil refineries," *Review of Economics and Statistics*, 83, 498–510.

Bertrand, M., E. Duflo, and S. Mullainathan (2004): "How Much Should We Trust Differences-In-Differences Estimates?," *The Quarterly Journal of Economics*, 119, 249–275.

Bloom, N., and J. Van Reenen (2007): "Measuring and Explaining Management Practices Across Firms and Countries," *The Quarterly Journal of Economics*, 122, 1351–1408.

Bosquet, B. (2000): "Environmental Tax Reform: Does It Work? A Survey of the Empirical Evidence," *Ecological Economics*, 34, 19–32.

Braguinsky, S., A. Ohyama, T. Okazaki, and C. Syverson (2015): "Acquisitions, Productivity, and Profitability: Evidence from the Japanese Cotton Spinning Industry," *The American Economic Review*, 105, 2086–2119.

Bredin, D., and C. Muckley (2011): "An Emerging Equilibrium in the EU Emissions Trading Scheme," *Energy Economics*, 33, 353–362.

Bushnell, J. B., H. Chong, and E. T. Mansur (2013): "Profiting from regulation: Evidence from the European carbon market," *American Economic Journal: Economic Policy*, 5, 78–106.

Calel, R., and A. Dechezleprêtre (2016): "Environmental policy and directed technological change: evidence from the European carbon market," *Review of Economics and Statistics*, 98.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014a): "Robust Data-Driven Inference in the Regression-Discontinuity Design," *The Stata Journal*, 14, 909–946.

——— (2014b): "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 82, 2295–2326.

Card, D., D. Lee, Z. Pei, and A. Weber (2012): "Inference on causal effects in a generalized regression kink design," *Econometrica*, 83, 2453–2483.

Chen, Y., W. K. Härdle, and U. Pigorsch (2010): "Localized Realized Volatility Modelling," *Journal of the American Statistical Association*, 105, 1376–1393.

Chevallier, J. (2009): "Carbon Futures and Macroeconomic Risk Factors: A View From the EU ETS," *Energy Economics*, 31, 614–625.

——— (2011a): "Evaluating the Carbon-Macroeconomy Relationship: Evidence from Threshold Vector Error-Correction and Markov-switching VAR Models," *Economic Modelling*, 28, 2634–2656.

——— (2011b): "A Model of Carbon Price Interactions with Macroeconomic and Energy Dynamics," *Energy Economics*, 33, 1295–1312.

Chevallier, J., F. Ielpo, and L. Mercier (2009): "Risk aversion and institutional information disclosure on the European carbon market: a case-study of the 2006 compliance event," *Energy Policy*, 37, 15–28.

Chevallier, J., Y. Le Pen, and B. Sévi (2011): "Options introduction and volatility in the EU ETS," *Resource and energy economics*, 33, 855–880.

Coase, R. H. (1960): "The Problem of Social Cost," *Journal of Law and Economics*, 3, 1–44.

Collard-Wexler, A., and J. De Loecker (2015): "Reallocation and Technology: Evidence from the US Steel Industry," *The American Economic Review*, 105, 131–171.

Commins, N., S. Lyons, M. Schiffbauer, and R. S. Tol (2011): "Climate policy and corporate behaviour," *Energy Journal*, 32, 51–63.

Conrad, C., D. Rittler, and W. Rotfuss (2012): "Modeling and explaining the dynamics of European Union Allowance prices at high-frequency," *Energy Economics*, 34, 316–326.

Creti, A., P. Jouvet, and V. Mignon (2012): "Carbon Price Drivers: Phase I versus Phase II Equilibrium?," *Energy Economics*, 34, 327–334.

Criscuolo, C., and R. Martin (2009): "Multinationals and U.S. Productivity Leadership: Evidence from Great Britain," *Review of Economics and Statistics*, 91, 263–281.

Crocker, T. D. (1966): "The structuring of atmospheric pollution control systems," in *The economic of air pollution*, ed. by H. Wolozin. Norton, New York.

Cropper, M. L., and W. E. Oates (1992): "Environmental economics: a survey," *Journal of economic literature*, 30, 675–740.

Dales, J. H. (1968): *Pollution, Property, and Prices*. University of Toronto Press, Toronto.

Daskalakis, G., D. Psychoyios, and R. Markellos (2009): "Modeling CO2 Emission Allowance Prices and Derivatives: Evidence from the European Trading Scheme," *Journal of Banking and Finance*, 33, 1230–1241.

De Loecker, J. (2007): "Do Exports Generate Higher Productivity? Evidence from Slovenia," *Journal of International Economics*, 73, 69–98.

——— (2013): "Detecting Learning by Exporting," *American Economic Journal: Microeconomics*, 5, 1–21.

De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2016): "Prices, Markups, and Trade Reform," *Econometrica*, 84, 445–510.

De Loecker, J., and F. Warzynski (2012): "Markups and Firm-Level Export Status," *The American Economic Review*, 102, 2437–2471.

Demailly, D., and P. Quirion (2008): "European Emission Trading Scheme and Competitiveness: A Case Study on the Iron and Steel Industry," *Energy Economics*, 30, 2009–2027.

Destatis (2015): "Deutscher Außenhandel: Export und Import im Zeichen der Globalisierung," https://www.destatis.de/DE/Publikationen/Thematisch/Aussenhandel/Gesamtentwicklung/AussenhandelWelthandel5510006139004.pdf, extracted on 29th of April, 2016.

Diebold, F. X., and A. Inoue (2001): "Long Memory and Regime Switching," *Journal of Econometrics*, 105, 131–159.

Doraszelski, U., and J. Jaumandreu (2013): "R& D and Productivity: Estimating Endogenous Productivity," *Review of Economic Studies*, 80, 1338–1383.

Ellerman, A. D., F. J. Convery, and C. De Perthuis (2010): *Pricing Carbon: The European Union Emissions Trading Scheme*. Cambridge University Press, Cambridge.

Ellerman, A. D., P. L. Joskow, R. Schmalensee, J.-P. Montero, and E. M. Bailey (2000): *Markets for Clean Air: the U.S. Acid Rain Program.* Cambridge University Press, Cambridge.

Engle, R. (2000): "The Econometrics of Ultra-High-Frequency Data," *Econometrica*, 68, 1–22.

European Commission (2011): "Proposal for a Directive on energy efficiency and repealing Directives 2004/8/EC and 2006/32/EC," [COM/2011/370].

European Council (2014): "Conclusion EUCO 189/14 of the European Council of 23 nd 24 October 2004," *Official Journal of the European Union*, CO EUR 13.

European Parliament and Council (2003): "Directive 2003/87/EC of the European Parliament and of the Council of 13 October 2003 establishing a scheme for greenhouse gas emission allowance trading within the Community and amending Council Directive 96/61/EC," *Official Journal of the European Union*, L 275, 32–46.

——— (2004): "Directive 2004/101/EC of the European Parliament and of the Council of 27 October 2004 amending Directive 2003/87/EC establishing a scheme for greenhouse gas emission allowance trading within the Community, in respect of the Kyoto Protocol's project mechanisms," *Official Journal of the European Union*, L 338, 18–30.

——— (2009): "Directive 2009/29/EC of the European Parliament and of the Council of 23 April 2009 amending Directive 2003/87/EC so as to improve and extend the greenhouse gas emission allowance trading scheme of the Community," *Official Journal of the European Union*, L 140, 63–87.

Eurostat (2014): "Electricity Prices of Industrial (nrg_pc_205_h) and Domestic Consumers (nrg_pc_204_h)," extracted on 4th of August, 2014.

——— (2016): "Greenhouse gas emission statistics," http://ec.europa.eu/eurostat/statistics-explained/index.php/Greenhouse_gas_emission_statistics, extracted on 29th of April, 2016.

Fabra, N., and M. Reguant (2014): "Pass-through of emissions costs in electricity markets," *The American Economic Review*, 104, 2872–2899.

FAMA, E. F., AND K. R. FRENCH (1987): "Commodity futures prices: some evidence on forecast power, premiums, and the theory of storage.," *The Journal of Business*, 25, 23–49.

——— (1989): "Business Conditions and Expected Returns on Stocks and Bonds," *Journal of Financial Economics*, 60, 55–73.

FEDERAL ENVIRONMENTAL AGENCY (2008): "Bestimmung spezifischer Treibhausgas-Emissionsfaktoren für Fernwärme," *Federal Environmental Agency - Climate Change*, 2008/08.

——— (2012): "Einheitliche Stoffwerte für Emissionsfaktoren, Heizwerte und Kohlenstoffgehalte für Brennstoffe, Rohstoffe und Produkte," https://www.dehst.de/SharedDocs/Downloads/Archiv/Zuteilung_2008-2012/Anhang01_Stoffliste.html, extracted on 29th of April, 2016.

——— (2014): "Entwicklung der spezifischen Kohlendioxid-Emissionen des deutschen Strommix in den Jahren 1990 bis 2013," *Federal Environmental Agency - Climate Change*, 2014/23.

FERRIS, A. E., R. J. SHADBEGIAN, AND A. WOLVERTON (2014): "The Effect of Environmental Regulation on Power Sector Employment: Phase I of the Title IV SO2 Trading Program," *Journal of the Association of Environmental and Resource Economists*, 1, 521–553.

FOWLIE, M., S. P. HOLLAND, AND E. T. MANSUR (2012): "What Do Emissions Markets Deliver and to Whom? Evidence from Southern California's NOx Trading Program," *American Economic Review*, 102, 965–93.

FOWLIE, M., AND J. M. PERLOFF (2013): "Distributing pollution rights in cap-and-trade programs: are outcomes independent of allocation?," *Review of Economics and Statistics*, 95, 1640–1652.

GERSTER, A. (2015): "Do Electricity Prices Matter? Plant-Level Evidence from German Manufacturing," Available at SSRN: https://ssrn.com/abstract=2603211.

Granger, C., and N. Hyung (2004): "Occasional Structural Breaks and Long Memory with an Application to the S&P500 Absolute Returns," *Journal of Empirical Finance*, 11, 399–421.

Gray, S. F. (1996): "Modeling the Conditional Distribution of Interest Rates as a Regime-switching Process," *Journal of Financial Economics*, 42, 27–62.

Gray, W. B., R. J. Shadbegian, C. Wang, and M. Meral (2014): "Do EPA regulations affect labor demand? Evidence from the pulp and paper industry," *Journal of Environmental Economics and Management*, 68, 188–202.

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 Manufactures," *Journal of Political Economy*, 110, 1175–1219.

Greenstone, M., J. A. List, and C. Syverson (2012): "The Effects of Environmental Regulation on the Competitiveness of U.S. Manufacturing," Available at NBER Working Paper Series: http://www.nber.org/papers/w18392.

Greestone, M. (2004): "Did the Clean Air Act Cause the Remarkable Decline in Sulfur Dioxide Concentrations?," *Journal of Environmental Economics and Management*, 47, 585–611.

Hahn, J., P. Todd, and W. van der Klaauw (2001): "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 69, 201–209.

Hall, B. H., and J. Mairesse (1995): "Exploring the relationship between R&D and productivity in French manufacturing firms," *Journal of Econometrics*, 65, 263–293.

Hamilton, J. D. (1989): "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57, 357–384.

Hanna, R. (2010): "US environmental regulation and FDI: evidence from a panel of US-based multinational firms," *American Economic Journal: Applied Economics*, 2, 158–189.

Harrison, A., B. Hyman, L. Martin, and S. Nataraj (2015): "When do firms go green? Comparing price incentives with command and control regulations in India," NBER Working Paper 21763: http://www.nber.org/papers/w21763.

Heckman, J., H. Ichimura, J. Smith, and P. Todd (1998): "Characterizing selection bias using experimental data," *Econometrica*, 66, 1017–1098.

Heckman, J. J., H. Ichimura, and P. E. Todd (1997): "Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme," *The Review of Economic Studies*, 64, 605–654.

Henderson, V. (1996): "Effects of Air Quality Regulation," *American Economic Review*, 86, 789–813.

Hicks, J. R. (1932): *The Theory of Wages*. Macmillan, New York.

Hintermann, B. (2010): "Allowance Price Drivers in the First Phase of the EU ETS," *Journal of Environmental Economics and Management*, 59, 43–56.

Hintermann, B., S. Peterson, and W. Rickels (2016): "Price and Market Behavior in Phase II of the EU ETS: A Review of the Literature," *Review of Environmental Economics and Policy*, 10, 108–128.

Holland, P. W. (1986): "Statistics and causal inference," *Journal of the American Statistical Association*, 81, 945–960.

Horowitz, J. L. (2001): "The Bootstrap," in *Handbook of Econometrics*, ed. by J. J. Heckman, and E. E. Learner, vol. 5, pp. 3159–3228. Elsevier Science, Oxford.

Imbens, G., and K. Kalyanaraman (2012): "Optimal Bandwidth Choice for the Regression Discontinuity Estimator," *The Review of Economic Studies*, 79, 933–959.

Imbens, G. W., and T. Lemieux (2008): "Regression Discontinuity Designs: A Guide to Practice," *Journal of Econometrics*, 142, 615–635.

Imbens, G. W., and J. M. Wooldridge (2009): "Recent developments in the econometrics of program evaluation," *Journal of Economic Literature*, 47, 5–86.

Jaraitė, J., and C. Di Maria (2016): "Did the EU ETS Make a Difference? An Empirical Assessment Using Lithuanian Firm-Level Data," *The Energy Journal*, 37, 1–23.

Kahn, M. E., and E. T. Mansur (2013): "Do local energy prices and regulation affect the geographic concentration of employment?," *Journal of Public Economics*, 101, 105–114.

Keppler, J. H., and M. Mansanet-Bataller (2010): "Causalities between CO2, Electricity, and Other Energy Variables During Phase I and Phase II of the EU ETS," *Energy Policy*, 38, 3329–3341.

Kim, C.-J. (1994): "Dynamic linear models with Markov-switching," *Journal of Econometrics*, 60, 1–22.

Klaassen, F. (2002): "Improving GARCH Volatility Forecasts with Regime-switching GARCH," *Empirical Economics*, 27, 363–394.

Klemetsen, M. E., K. E. Rosendahl, and A. L. Jakobsen (2016): "The impacts of the EU ETS on Norwegian plants' environmental and economic performance," NMBU Working Papers 3/2016. Available at https://www.nmbu.no/sites/default/files/pdfattachments/hh_wp_3_2016_0.pdf.

Klepper, G., and S. Peterson (2004): "The EU Emissions Trading Scheme, Allowance Prices, Trade Flows, Competitiveness Effects," *European Environment*, 14, 201–218.

Knittel, C. R. (2002): "Alternative Regulatory Methods and Firm Efficiency: Stochastic Frontier Evidence from the U.S. Electricity Industry," *Review of Economics and Statistics*, 84, 530–540.

Koch, N. (2014): "Dynamic linkages among carbon, energy and financial markets: a smooth transition approach," *Applied Economics*, 46, 715–729.

Koch, N., S. Fuss, G. Grosjean, and O. Edenhofer (2014): "Causes of the EU ETS price drop: Recession, CDM, renewable policies or a bit of everything? - New evidence," *Energy Policy*, 73, 676–685.

Koch, N., G. Grosjean, S. Fuss, and O. Edenhofer (2016): "Politics matters: regulatory events as catalysts for price formation under cap-and-trade," *Journal of Environmental Economics and Management*, 78, 121–139.

Lee, D. S. (2008): "Randomized Experiments from Non-random Selection in U.S. House Elections," *Journal of Econometrics*, 142, 675–697.

Lee, D. S., and T. Lemieux (2010): "Regression Discontinuity Designs in Economics," *Journal of Economic Literature*, 48, 281–355.

Levinsohn, J., and A. Petrin (2003): "Estimating Production Functions Using Inputs to Control for Unobservables," *The Review of Economic Studies*, 70, 317–341.

Linn, J. (2010): "The effect of cap-and-trade programs on firms' profits: Evidence from the Nitrogen Oxides Budget Trading Program," *Journal of Environmental Economics and management*, 59, 1–14.

List, J. A., D. L. Millimet, P. G. Fredriksson, and W. W. McHone (2003): "Effects of environmental regulations on manufacturing plant births: evidence from a propensity score matching estimator," *Review of Economics and Statistics*, 85, 944–952.

Lutz, B. J. (2016): "Emissions Trading and Productivity: Firm-level Evidence from German Manufacturing," ZEW Discussion Paper No. 16-067.

Mansanet-Bataller, M., A. Pardo, and E. Valor (2007): "CO2 Prices, Energy and Weather," *Energy Journal*, 28, 67–86.

Marschak, J., and W. H. Andrews (1944): "Random Simultaneous Equations and the Theory of Production," *Econometrica*, 12, 143–205.

Martin, R., L. B. de Preux, and U. J. Wagner (2014): "The Impact of a Carbon Tax on Manufacturing: Evidence from Microdata," *Journal of Public Economics*, 117, 1–14.

Martin, R., M. Muûls, and U. J. Wagner (2016): "The Impact of the EU ETS on Regulated Firms: What is the Evidence After Ten Years?," *Review of Environmental Economics and Policy*, 10, 129–148.

McCrary, J. (2008): "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test," *Journal of Econometrics*, 142, 698 – 714, The regression discontinuity design: Theory and applications.

Medina, V., and A. Pardo (2012): "Is the EUA a New Asset Class?," *Quantitative Finance*, 13, 637–653.

Medina, V., Á. Pardo, and R. Pascual (2014): "The timeline of trading frictions in the European carbon market," *Energy Economics*, 42, 378–394.

Meeusen, W., and J. Van den Broeck (1977): "Efficiency estimation from Cobb-Douglas production functions with composed error," *International Economic Review*, 18, 435–444.

Mikosch, T., and C. Stărică (2004): "Non-stationarities in Financial Time Series, the Long Range Dependence and the IGARCH Effects," *Review of Economics and Statistics*, 86, 378–390.

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

Milne, J. E. (2011): "Environmental Taxation in the United States: The Long View," *Lewis and Clark Law Review*, 15, 417–229.

Mizrach, B., and Y. Otsubo (2014): "The market microstructure of the European climate exchange," *Journal of Banking & Finance*, 39, 107–116.

Montagnoli, A., and F. de Vries (2010): "Carbon Trading Thickness and Market Efficiency," *Energy Economics*, 32, 1331–1336.

Montgomery, W. D. (1972): "Markets in Licenses and Efficient Pollution Control Programs," *Journal of Economic Theory*, 5, 395–418.

Oberndorfer, U., and K. Rennings (2007): "Costs and Competitiveness Effects of the European Union Emissions Trading Scheme," *European Environment*, 17, 1–17.

Olley, G. S., and A. Pakes (1996): "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64, 1263–1297.

PAOLELLA, M., AND L. TASCHINI (2008): "An Econometric Analysis of Emission Allowances Prices," *Journal of Banking and Finance*, 32, 2022–2032.

PEARL, J. (2000): *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York.

PERI, M., AND L. BALDI (2011): "Nonlinear Price Dynamics between CO2 Futures and Brent," *Applied Economics Letters*, 18, 1207–1211.

PETRICK, S., K. REHDANZ, AND U. J. WAGNER (2011): "Energy use patterns in German industry: Evidence from-level data," *Jahrbücher für Nationalökonomie und Statistik*, pp. 379–414.

PETRICK, S., AND U. J. WAGNER (2014): "The Impact of Carbon Trading on Industry: Evidence from German Manufacturing Firms," Available at SSRN: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2389800.

PIGOU, A. C. (1920): *The Economic of Welfare*. Macmillan and Company, London.

PLENK, J., G. NERB, K. WOHLRABE, AND T. O. BERG (2013): "CESifo World Economic Survey February 2013," *CESifo World Economic Survey*, 12, 1–25.

PORTER, J. (2003): "Estimation in the Regression Discontinuity Model," Working Paper, University of Wisconsin. Available at http://www.ssc.wisc.edu/∼jrporter/reg_discont_2003.pdf.

PORTER, M. E. (1991): "America's green strategy.," *Scientific American*, 264, 168.

PORTER, M. E., AND C. VAN DER LINDE (1995): "Toward a new conception of the environment competitiveness relationship.," *Journal of Economic Perspectives*, 9, 97–118.

REBOREDO, J. C. (2013): "Modeling EU allowances and oil market interdependence. Implications for portfolio management," *Energy Economics*, 36, 471–480.

RESEARCH DATA CENTRES OF THE STATISTICAL OFFICES GERMANY (2014): "AFiD-Panel Industrial Units, AFiD-Module Use of Energy, and Cost Structure Survey 1995-2012," Further information: http://www.forschungsdatenzentrum.de/en/index.asp.

RITTLER, D. (2012): "Price discovery and volatility spillovers in the European Union emissions trading scheme: A high-frequency analysis," *Journal of Banking & Finance*, 36, 774–785.

ROBERTS, M. J., AND M. SPENCE (1976): "Effluent charges and licenses under uncertainty," *Journal of Public Economics*, 5, 193–208.

ROBLES-FERNANDEZ, M. D., L. NIETO, AND M. A. FERNANDEZ (2004): "Nonlinear Intraday Dynamics in Eurostoxx50 Index Markets," *Studies in Nonlinear Dynamics& Econometrics*, 8, Article 3.

ROTFUSS, W. (2011): "On the High-Frequency Price Reactions of European Union Allowances to News," Ph.D. thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg.

RUBIN, D. (1974): "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies," *Journal of Educational Psychology*, 66, 688–701.

——— (1977): "Assignment to a Treatment Group on the Basis of a Covariate," *Journal of Educational Statistics*, 2, 1–26.

SADORSKY, P. (2002): "Time-varying Risk Premiums in Petroleum Futures Prices,," *Energy Economics*, 24, 539–556.

SMITH, H. L. (1997): "Matching With Multiple Controls to Estimate Treatment Effects in Observational Studies," *Sociological Methodology*, 27, 325–353.

SYVERSON, C. (2011): "What Determines Productivity?," *Journal of Economic Literature*, 49, 326–365.

THE FEDERATION OF GERMAN INDUSTRIES (2008): "Kostenfaktor Strom: Wettbewerbsfähige Preise für die Industrie - Cost Factor Electricity: Competitive Prices for the Industry," Available at http://www.bdi.eu/publikationen-flyer_1873.htm extracted on the 8th of June 2015. .

——— (2016): "CO2-Preis setzt Wettbewerbsfähigkeit europäischer Industrien zusätzlich unter Druck - CO2 price puts competitiveness of European industry under pressure," Available at http://bdi.eu/artikel/news/co2-preis-setztwettbewerbsfaehigkeit-europaeischer-industrien-zusaetzlich-unter-druck/ extracted on the 12th of January 2016.

van Biesebroeck, J. (2003): "Productivity Dynamics with Technology Choice: An Application to Automobile Assembly," *The Review of Economic Studies*, 70, 167–198.

Wagner, U. J., M. Muûls, R. Martin, and J. Colmer (2014): "The Causal Effects of the European Union Emissions Trading Scheme: Evidence from French Manufacturing Plants," Mimeo.

Walker, W. R. (2011): "Environmental regulation and labor reallocation: Evidence from the Clean Air Act," *The American Economic Review*, 101, 442–447.

——— (2013): "The transitional costs of sectoral reallocation: evidence from the clean air act and the workforce," *Quaterly Journal of Economics*, 128, 1787–1835.

Weitzman, M. L. (1974): "Prices vs. quantities," *The review of economic studies*, 41, 477–491.

World Bank (2012): *State and Trends of the Carbon Market 2012*. Carbon Finance at the World Bank.