Has emissions trading reduced emissions from Norwegian firms?

An empirical analysis of Norwegian manufacturing industry

Maren Holthe Hedne

Dissertation supervisor: Stephen Smith

Dissertation submitted in part-fulfilment of the MSc in Economics
University College London
September 2020

Abstract

The European Union Emissions Trading System (EU ETS) has been the cornerstone climate policy of the EU since 2005. It regulates around half of EU greenhouse gas emissions, yet its actual impact – nationally and across the Union – is uncertain. Applying Norwegian administrative plant-level data from 2001 to 2019, this paper contributes new and updated evidence of the impact on Norwegian ETS firms, compared to the counterfactual scenario of non-ETS firms. I estimate an impact at around 20 % lower emissions in phase II (2008–2012). Surprisingly, I estimate around 30 % higher emissions in phase III (2013–2019). While estimates vary somewhat across model specifications, the sign and relative magnitude are consistent. I present a brief discussion of the findings, and leave some recommendations for future research.

Acknowledgements

This dissertation would not have come about without the support of some key people. The guidance of my supervisor Stephen Smith inspired and helped structure and frame the project. My co-supervisor Elisabeth Thuestad Isaksen at the Ragnar Frisch Centre for Economic Research in Oslo has guided me through the empirical challenges of the process and provided invaluable advice. The office space at the Frisch Centre this summer has provided a valuable community of economists in a time of isolation for us all, and has inspired me to continue my research as a PhD student.

My partner Håvard Sagbakken Saanum has lent unwavering moral support and encouragement, as always. To my closest allies among the students at the MSc – your hard work, insights and ideas have given me the motivation and energy to finish the project. Thank you all.

Contents

List of figures

List of tables

1 Introduction

The EU emissions trading system (EU ETS) covers all 28 EU countries and the three EEA countries Norway, Liechtenstein and Iceland. By regulating 11 000 installations and aviation, the EU ETS covers approximately 45 % of greenhouse gas emissions within the participating countries (Commission 2020). What started off as the first large-scale greenhouse gas emissions trading system soon became the cornerstone climate policy of the EU. Through emissions trading the EU would deliver cost-efficient abatement without imposing drastic costs or risks of carbon leakage on its industry. Around one half of Norwegian greenhouse gas (GHG) emissions are covered by the EU ETS. However, since phase III began in 2013, around 90 % of Norwegian industrial emissions are covered. This raises a need to understand how manufacturing firms are impacted by EU ETS regulation and whether it has succeeded in delivering abatement in Norway. The jury is still out on whether it has succeeded thus far in delivering on its highly ambitious goals. This paper aims to contribute some evidence to that debate.

This paper is conceptually inspired by Klemetsen, Rosendahl, and Jakobsen (2020). I apply difference-in-differences estimators with and without inverse propensity score weighting to identify the causal impact of the EU ETS on Norwegian ETS firms compared to non-ETS firms, using plant-level data from 2001 to 2019. My main contribution is providing empirical evidence and analysis which extends to phase III. The main findings in 2008–2013 are in line with those of Klemetsen, Rosendahl, and Jakobsen (2020). Under some strict assumptions I estimate reduced emissions of around 20–40 % in phase II (2008–2012), and increased emissions of around 30–60 % in phase III (2013–2019). I discuss the findings along two dimensions – first, whether this can be taken to mean that the EU ETS has been inefficient in delivering abatement, and second, whether it is better explained by deviating marginal abatement costs across the EU.

The theoretical and historical backdrop – a simple theory of abatement under emissions trading, the phases of the EU ETS and the Norwegian regulatory context – are presented in chapter 2. The recent literature on the EU ETS is summarised in chapter 3, and the data are described in chapter 4. Chapter 5 discusses the methodologies applied in the empirical analysis. Chapter 6 presents the results. Chapter 7 discusses these the main findings and some implications, and chapter 8 contains some concluding remarks.

2 Background and theoretical considerations

2.1 Abatement under emissions trading

The goal of emissions trading is mitigation of greenhouse gases as an externality by cost-minimising means (Martin, Muûls, and Wagner 2016). As climate change is a global problem, the socially optimal way of mitigating climate change is to reduce emissions where the marginal abatement cost is lowest, regardless of geographical location. The EU ETS sets a cap on emissions, and ETS firms have to hold a European Union allowance (EUA) for each tonne of CO_2 equivalents they emit. The price of EUAs provides firms with incentives to reduce their emissions, while the cap ensures that the target level of GHG emissions is met.

Appendix describes the theory behind emissions trading in more detail. The cap should be set at a level that ensures the socially optimal level of pollution occurs. The allowances can be distributed either for free ('grandfathering') – based on output, technology or past emissions – or auctioned off. Auctioning is often viewed as advantageous because it ensures governmental income and can be used to reduce other, distortionary taxes (Goulder and Parry 2008). As I show below, however, the allowance price depends on the marginal abatement cost (MAC) of firms rather than choices of allocation. It is also the MAC and allowance price that determines the chosen level of abatement for each firm, not the total cost pollution. Economic theory does therefore not predict any difference in abatement levels for grandfathering and auctioning schemes. Finally, in the EU, member-state control mechanisms ensure that firms comply with regulation (Wettestad and Jevnaker 2018). Overall, this means that the EU ETS should deliver the same abatement independent of the allocation schemes, as long as the cap has been tight and compliance has been achieved. The EU ETS does not, however, restrict where pollution occurs geographically. Instead, this is determined by the allowance market.

Firms can adapt to emissions trading by reducing their emissions (either by reducing their production or emission intensity) or participate in the market for allowances. In a simple, static model of emissions trading I denote the cap on emissions by κ. The required abatement by the economy as a whole, $\mathcal{E}$, is then given by the difference between business-as-usual emissions E and the cap:

$$\mathcal{E} = E - \kappa \tag{2.1}$$

Assuming the difference between business-as-usual emissions and the cap is positive, the drivers of the allowance price are the magnitude of $\mathcal{E}$ and the MAC of participating agents (Hintermann, Peterson, and Rickels 2016). Suppose the economy consists of only two firms, firm 1 and 2, which differ in their MAC [1]. Assuming perfect compliance, the sum of emissions by firm 1 and firm 2 therefore cannot exceed the cap. Firms are assumed to have a positive and increasing MAC schedule, such that the cap will always be exhausted. In figure 2.1, firm 1 has a higher MAC than firm 2. If both firms are profit maximising, they will reduce their emissions until the MAC is higher than the market price of allowances. The abatement of each firm $i = \{1, 2\}$ is denoted by $\mathcal{E}_i \equiv E_i - \kappa_i$, where κ_i denotes the post-abatement emissions of firm i. Both firms are willing to sell allowances at any price that at least equals their MAC of one extra unit. The equilibrium is therefore given by the point where the two MAC schedules intersect, $(\mathcal{E}_1^*, \mathcal{E}_2^*)$, and the allowance price P. The allowance price will be given by the MAC of firm 2 of abating one more unit.

Figure 2.1 illustrates that firms under cap-and-trade behave in a way that achieves cost-minimising abatement. The model further illustrates that where abatement occurs depends on the MAC of the firms. The firm with the lowest MAC schedule, here firm 2, conducts more abatement than firm 1.

The MAC of a firm depends on availability of technology, input substitutes and to what extent abatement has occurred in the past. A large share of hydro energy means that the Norwegian power market is characterised by more than 95 % renewable energy and very low electricity prices. Consequently, Norwegian power-intensive industries have to a large extent based their production on 'clean' inputs (NEA 2010). Further, EU and EEA member-states have had varying levels of ambition in their climate policies prior to the ETS. MACs are generally assumed to be increasing in the level of abatement, this might

1. The limitation to two firms is only a simplifying assumption. The overall conclusion extends to any number of firms.

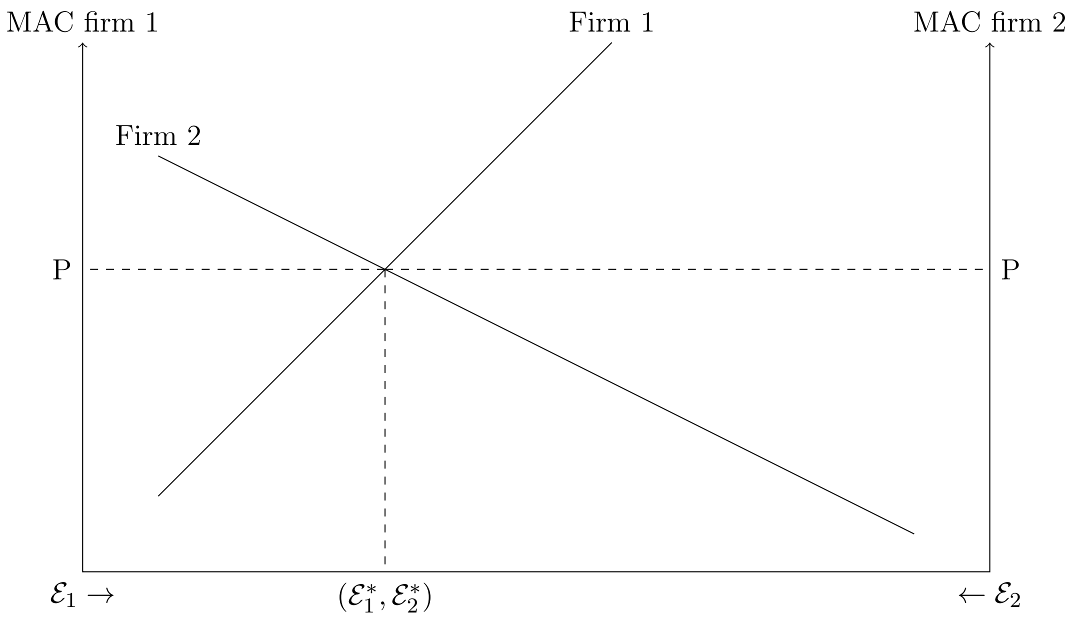

Figure 2.1: Abatement under cap-and-trade

Note: In an economy with only two firms, $\mathcal{E}$ denotes the total cap on allowances and is represented by the entire horizontal line. Each firm $i = \{1, 2\}$ conducts abatement given by $\mathcal{E}_i$. The equilibrium price P and abatement distribution $(\mathcal{E}_1^*, \mathcal{E}_2^*)$ is determined by the point where the marginal abatement cost (MAC) curves intersect.

suggest that Norwegian industry has higher MACs than similar industries in Europe.

The simple model described here further implies that low allowance prices does not in itself mean that there is no impact on emissions. Arguably, firms could initially be harvesting 'low-hanging fruit' such that initial MACs are very low. It is, however, a concern if allowance prices stay well below the social price of carbon long-term, as this means that the optimal equilibrium is likely not to be reached. This efficiency concern is part of what calls for empirical analyses of the EU ETS – in light of the low allowance prices, to what extent and where has abatement occurred?

2.2 Phases and empirical assessment of the EU ETS

The discussion of an EU ETS began following the Kyoto protocol. The European Commission had previously attempted to launch a tax on carbon, but proved unsuccessful (Wettestad and Jevnaker 2018). The first ETS directive (Directive 2003/87/EC) was adopted in 2003 and set out the regulation for phases I and II. The main components of the EU ETS have remained constant over time, but there has been a clear trend of increased centralisation and 'tightening' of the cap. The history and background of the EU ETS and emissions trading in Norway is described in detail in appendix and summarised in figure 2.3.

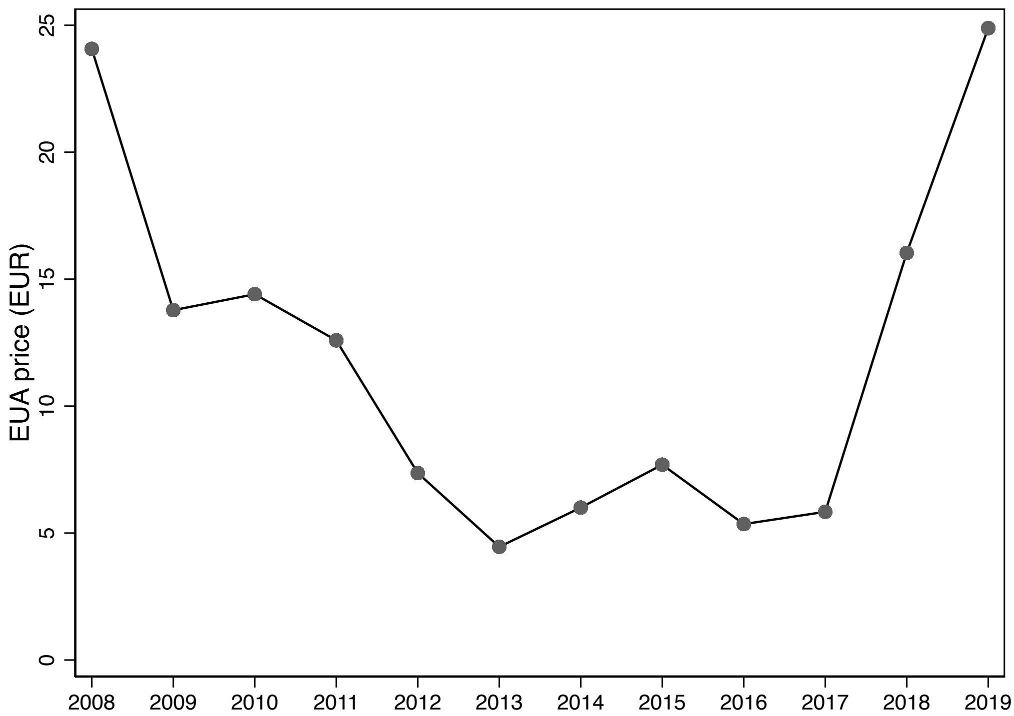

Figure 2.2: Yearly mean price of EUAs, 2008*–2019

Note: The figure plots the yearly mean of the daily price of EU allowances from 2008–2019. The data are collected from Ember (formerly Sandbag) on 30.07.2020.
***** The observations start on 07.04.2008, due to data availability.

The EU ETS set off with a pilot phase (2005–2007), where the covered industries (mainly power generation and energy-intensive industries) and criteria were set centrally, but the cap and implementation was determined by each member-state. While Norway was not a member of the EU ETS in the first phase, it had introduced its own emissions trading system with the aim of integrating with the EU ETS from phase II. Norwegian allowances were distributed for free based on emissions in 1998–2001.

Norway joined the EU ETS from the Kyoto phase (phase II, 2008–2012). This is also the first phase covered by my analysis. Allowances were mainly distributed for free, and could be 'banked' for future years. There was no central cap, but member-states determined their own allocation rules. In phase II, Norway distributed allowances equivalent to 80 % of plant emissions in 2005. The EU ETS also allowed some offsets of CDM (Clean Development Mechanism) and JI (Joint Initiative) allowances. EUA prices started off relatively high but soon declined when the 2008/2009 recession coincided with the realisation that there was an over-allocation of allowances. In general, however, econometric studies have found higher estimates of abatement in this phase than in phase I. Klemetsen, Rosendahl, and Jakobsen (2020) estimate that phase II led to around 30 % lower emissions in Norwegian ETS firms, while a working paper by Petrick and Wagner (2014) estimate abatement at about one-fifth for German firms in the first half of phase II.

Phase III (Directive 2008/101/EC, Directive 2009/29/EC) entailed the first major reform of the EU ETS, and is more centralised in design and more comprehensive in scope (Wettestad and Jevnaker 2018). According to harmonised allocation rules, an increasing share of allowances are now auctioned, and only those firms deemed most at risk of relocation receive some allowances for free. Grandfathering is determined by a 'benchmarking' rule which in addition to carbon leakage exposure is largely output-based (Union 2015), which may act as an implicit production subsidy and can hence affect production decisions (Rosendahl and Storrøsten 2011). A central cap is in place, with a linear reduction in allowances of 1.74 % per year. Additional industries were included, and the scope expanded from CO_2 emissions only to including N_2O and PFCs. In Norway, this primarily led to PFCs from one of the country's largest industries, aluminium production, also being covered.

The level of ambition of the EU ETS has been increasing over time. While the European Commission has long been aiming for a centralised system, some member-states and industry stakeholders have held back (Wettestad and Jevnaker 2018). The same trend is expected to continue in the coming phase IV (2021–2030). Following a large surplus in allowances flooding the market, a market stability reserve (MRS) has been established from 2019. The MRS withdraws allowances from the market if the surplus grows too large, and introduces additional allowances if supply becomes too low. This means that abatement is no longer limited to the targets set by the EU, but can in fact exceed them. This opens the door for more member-state policies aiming towards EU ETS ETS firms. Overall, the increased harmonisation and additional measures are reasons to expect a higher impact on emissions in the

EU of phase III than phase II.

The EU ETS has been surrounded by controversy since the beginning – environmentalists have strived towards a more ambitious programme, while industry representatives have been concerned about their competitiveness in a global market. As so often is the case in the EU, the system has emerged as a compromise between strong interest groups. Thus, firms may not have been able to draw clear conclusions about how the regulation would impact them until after its adoption. On the other hand, regulation has generally been adopted well in advance of implementation, as shown in figure 2.3. This means that anticipation effects are likely to be present, but perhaps not very long before the regulation was adopted. This has implications for the comparison between ETS firms and non-ETS firms and is discussed further in section 6.1.

2.3 The Norwegian regulatory context

For the last two decades, Norwegian land-based industry has been far from unregulated. This section summarises the context of climate policies in which Norwegian industry firms have made their decision on production and generation of GHG emissions, providing a useful backdrop for the analysis of the difference between ETS and non-ETS firms.

The Norwegian CO_2 tax has been in place since 1991 and constitutes the main instrument addressing emissions from the manufacturing industry. The tax regulates GHG emissions from combustion of fossil fuels (petrol, mineral oil, natural gas and LPG), mineral products and the petroleum industry. However, some energy-intensive industries are exempt. The petroleum sector faces both the EU ETS and a CO_2 tax on its emissions, and identification of the individual effect of the two measures on abatement is therefore infeasible. The petroleum sector is therefore left out of this analysis. Land-based industry emissions covered by the CO_2 tax have been exempt from the ETS since phase I (Klemetsen, Rosendahl, and Jakobsen 2020). For non-ETS firms the tax has remained in place, generally at higher nominal levels than the price of EUAs. However, not all emissions or industrial processes are covered. Beyond pointing out that non-ETS firms have also had incentives for abatement, I therefore refrain from making any statement about the relative carbon prices of the EU ETS and the CO_2 tax.

Klemetsen, Rosendahl, and Jakobsen (2020) control for the CO_2 tax through including a variable for plant-specific relative energy prices. I argue that such a measure would be endogenous because energy prices are likely to be impacted by emissions trading. Energy producers have been suspected of passing on the cost of EUAs to consumers (Hintermann, Peterson, and Rickels 2016). This is likely to have lead to an increase in electricity prices. While near all of Norway's electricity generation is renewable, the Nordic power market is increasingly tied to the power market of the European continent. Thus, electricity prices are one of the channels through which the ETS impacts emissions of firms, and conditioning on (relative) energy prices could therefore confound the results.

Beyond the CO_2 tax, there have been strategies in place to reduce emissions from the non-ETS sectors (see e.g. NEA 2010). As Klemetsen, Rosendahl, and Jakobsen (2020) points out, some of these measures are industry-specific and are therefore not expected to apply differently to the control and treatment groups. Further, industry-specific measures are captured by conditioning on industry-fixed effects in all models in this paper. Without covering these measures or their costs specifically, it should be considered a possibility for the further analysis that the sum of measures aimed at non-ETS firms have had significant impact on emissions.

Overall, I believe that other regulatory instruments ought to be considered part of the counterfactual scenario. I therefore choose not to control for other climate policies, based on the assumption that the same measures would have been imposed on ETS firms in the absence of the EU ETS. While this assumption cannot be tested, multiple Norwegian governments have shown little hesitation over time to imposing a cost of emissions on the manufacturing industry[2], and it therefore seems likely that the measures aimed at non-ETS firms would have applied to all firms in the absence of ETS regulation. This means that the results of my analysis should be interpreted as changes in emissions from ETS firms, compared to firms that do not face EU ETS regulation – but possibly other sources of regulation.

2. For instance, the government lobbied for all permits to become auctioned from phase III and made it clear as early as in 2007 that Norwegian firms could no longer expect a grandfathering of permits after 2012 (https://stortinget.no/no/Saker-og-publikasjoner/Publikasjoner/Innstillinger/Odelstinget/2006-2007/inno-200607-100/1/#a2).

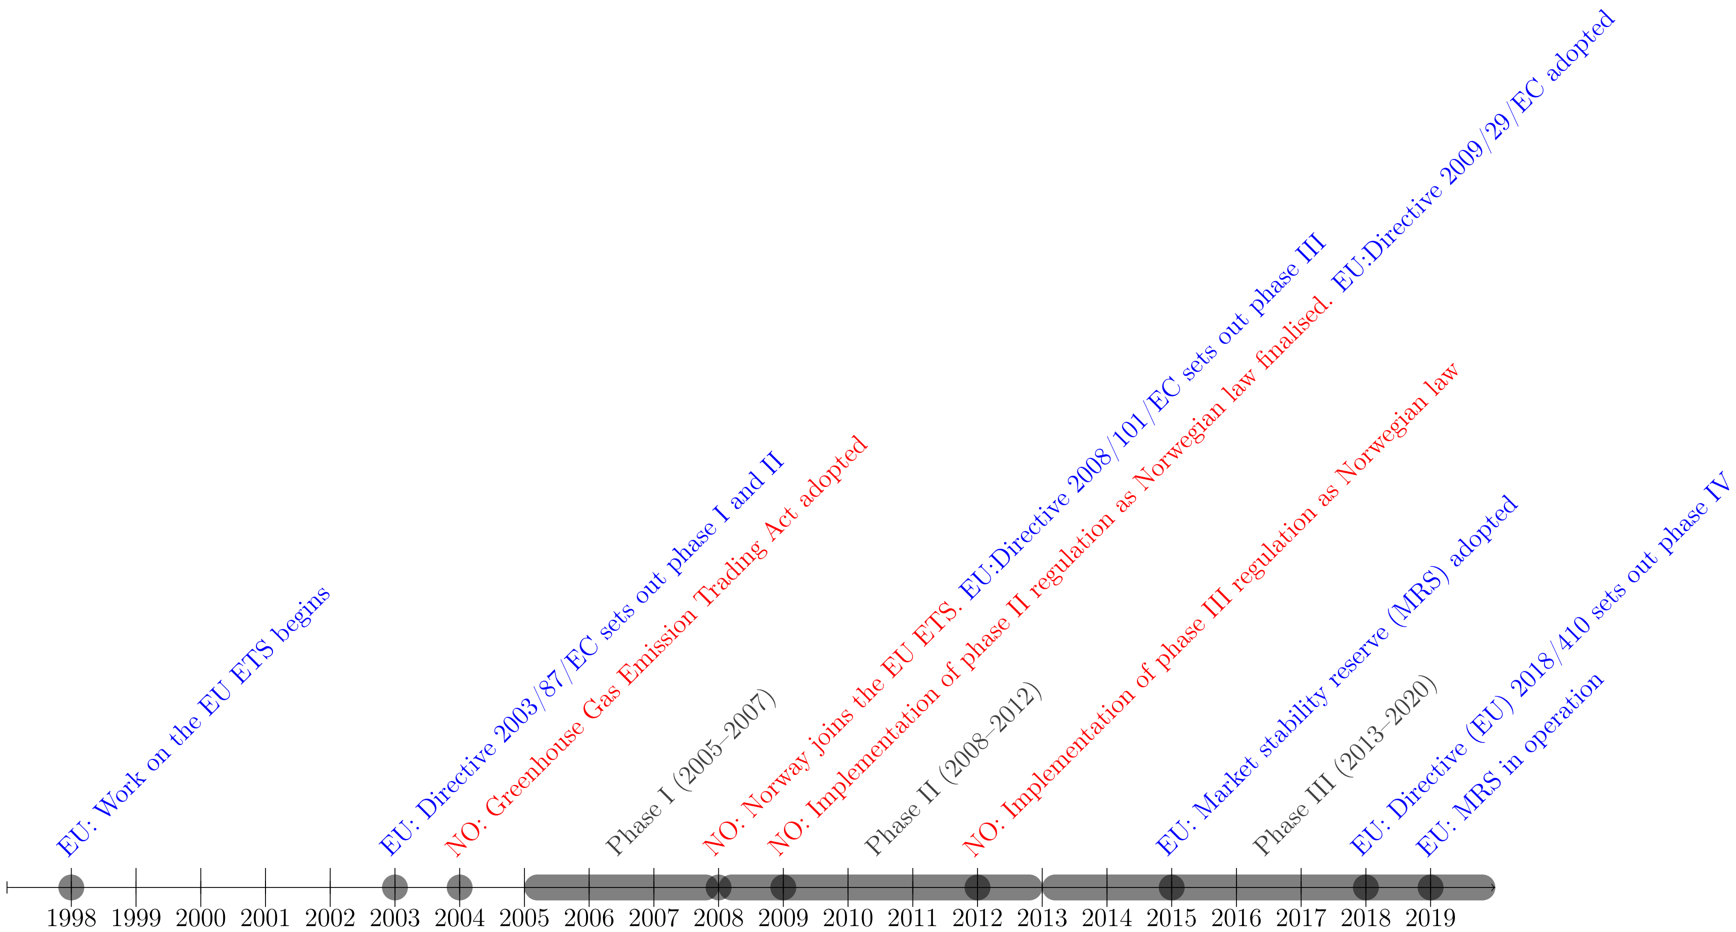

Figure 2.3: EU ETS timeline

Blue denotes EU-level events. Red denotes Norwegian events.

Note: The timeline describes the historical development of the Norwegian emissions trading system, its association and eventual inclusion in the EU ETS and key dates for emissions regulation. The Greenhouse Gas Emission Trading Act (2004) is the legal foundation of Norwegian participation in the EU ETS. See appendix for details. Only dates relevant for industrial plants are included.

3 Literature review

A range of research focusing on the economic impact of the EU ETS has been conducted. In this literature review I focus on recent causal research on the EU ETS conducted using firm or plant level data. Analysis using aggregate data, simulations and theoretical assessments are not covered here. [1] I emphasise their methodological and identification strategies to create a backdrop for the analytical sections of this paper.

Klemetsen, Rosendahl, and Jakobsen (2020)

Methodologically, this paper builds on Klemetsen, Rosendahl, and Jakobsen (2020). The authors use Norwegian plant-level data to assess the impact of the EU ETS on the economic and environmental performance of ETS firms. The authors use the same data as this paper on emissions from ETS firms and firms covered by the Norwegian Pollution Control Act from 2001–2013. They combine emissions data with data from Statistics Norway on firm-level economic variables. They apply DiD with propensity score matching, as well as a linear panel data fixed effects regression model. The time frame 2001–2013 allows them to assess phase I, phase II and the first year of phase III.

The dependent variables are GHG emissions, emissions intensity (measured as emissions relative to man hours), value added and labour productivity (measured as value added relative to man hours). As controls they include the relative energy price of fossil fuels to electricity, plant fixed effects (in the regression model), industry fixed effects, lagged number of employees and phase fixed effects. Finally, lagged employees serve as a measure of firm size.

Klemetsen, Rosendahl, and Jakobsen (2020) find a negative effect on emissions in phase II of about about 30 % (at the 10 % significance level). In other phases estimates are near zero and insignificant. No significant effect

1. For extensive reviews of previous and recent literature, see Martin et al. (2014) and Teixidó, Verde, and Nicolli (2019).

was found for emissions intensity. Interestingly, the authors estimate that EU ETS regulation is associated with increased economic performance of about 25 % across in phase II, but not phase I and III. Due to a lack of sufficient data the economic results are not reproduced here, and the findings on emissions receive more attention.

Other research

Bel and Joseph (2015) attempt to disentangle the impacts of the 2008/2009 recession and the EU ETS by applying a dynamic panel data model (as proposed by Arellano and Bond 1991). They apply four different specifications of either energy consumption from different sources or energy prices. By alternating between a dummy for ETS regulation and a variable for GDP growth in both specifications they assess which variable has more explanatory power. This paper differs from the others covered by this literature review in that it does not mainly estimate the treatment effects, but rather whether the ETS or recession were more powerful in explaining the changes in GHG emissions of ETS firms. Bel and Joseph (2015) criticise the majority of early papers (see for instance Ellerman and Buchner 2008) for being based on business-as-usual counterfactuals rather than application of firm-level data. Their main contribution is explicitly modelling the impact on emissions of the 2008/2009 shock to the European economy. More recent research has been based on firm-level data, an approach this paper shares.

A much-cited working paper by Petrick and Wagner (2014) is interested in the average treatment effect (ATT) on German manufacturing firms. They apply a combination of nearest neighbour and propensity score matching to a difference-in-difference regression weighed by the propensity score. Petrick and Wagner (2014) found that emissions from German manufacturing firms were reduced by one fifth in phase II, but no significant impact in phase I. Petrick and Wagner (2014), Klemetsen, Rosendahl, and Jakobsen (2020), and Löschel, Lutz, and Managi (2019) all deal implicitly with the 2008/2009 financial crisis by using difference-in-differences and assuming that both ETS and non-ETS firms were affected in the same manner by the recession.

Forbes and Zampelli (2019) also compare the EU ETS to another measure, namely the increasing wind energy penetration in the Irish energy mix. Using a time-series model they estimate the impact of the EU ETS and wind penetration on emissions from Irish electricity generation, using half-hourly data from 2015 to 2018. They find significant impact of the EU ETS, estimating that emissions would have been 6 % higher over the period without the EU ETS. However, wind energy penetration has a higher impact, and they estimate a substantial impact of the permit price on emissions.

Summary of research

In sum, this paper contributes to the growing literature on the causal impact of the EU ETS. This chapter and table 3.1 summarise the main studies conducted until now. Building on programme evaluation literature, the studies generally apply DiD estimation and some panel data regression. Published EU ETS evaluations typically focus on phase I and II. This likely reflects a lag in data publishing, access and publication of research. Klemetsen, Rosendahl, and Jakobsen (2020) and Forbes and Zampelli (2019) cover a few years of phase III, but neither provide a thorough evaluation of the phase nor a comparison with previous phases. I am aware of several forthcoming empirical studies on the EU ETS using microdata[2]. This shows that the topic is considered highly relevant, and that more causal evidence is likely to become available before long. As Teixidó, Verde, and Nicolli (2019) point out, the main gap in the literature is research on phase III. This will therefore be the main focus of this paper, with phase II receiving attention mostly for comparison with previous research.

2. See e.g. Wagner et al. (2020)

Paper	Methods	Industries	Years	Phases	Dependent variables	Key findings
Klemetsen, Rosendahl, and Jakobsen (2020)	DiD with matching + linear panel data regression	Manufacturing (Norway)	2001–2013	1 + 2 + 3	Emissions, emissions intensity, labour productivity, value added	No significant impacts in phase I. Weak evidence for abatement (around 30 %) and improved economic performance in phase II.
Petrick and Wagner (2014) [Working paper]	DiD with matching	Manufacturing (Germany)	2005–2010	1+2	Emissions, employment, emissions intensity	No significant impact on emissions in phase I, one-fifth in phase II (first half).
Bel and Joseph (2015)	Dynamic linear panel data regression	All (EU25*)	2005–2012	1+2	Emissions	Recession the main driver of emission reduction.
Löschel, Lutz, and Managi (2019)	DiD & stochastic frontier analysis	Manufacturing (Germany)	2003–2012	1+2	Economic performance, emissions	Phase I: Economic performance of ETS firms increases relative to non-ETS firms. Phase II: No significance.
Forbes and Zampelli (2019)	Time series regression	Electricity generation (Ireland)	2015–2018	3	Carbon emissions	Emissions would have been 6 % higher without the EU ETS.

Table 3.1: Recent literature

Note: Summarises some key points about the recent empirical EU ETS literature.
*I.e. the 25 EU member-states before 2007 – the current 27, the UK, and excluding the EEA, Bulgaria, Romania and Croatia.

4 Data

4.1 Data source

The analysis in this dissertation builds primarily on administrative data from the Norwegian Environmental Agency (Miljødirektoratet, NEA). The data consists of reported emissions from all Norwegian firms that are either covered by the EU ETS or the Norwegian Pollution Control Act[1]. By merging these data with a dataset of all EU ETS firms I obtain a dummy variable for EU ETS regulation. The datasets contain information on GHG emissions measured in tonnes of CO_2 equivalents. For the sake of comparison with other studies (particularly Klemetsen, Rosendahl, and Jakobsen 2020) I focus on plants belonging to industries B and C in SIC 2007, i.e. firms with two-digit NACE codes from 05 to 33 (except 06, petroleum and gas extraction). This includes mining, quarrying and manufacturing industries and excludes, among others, agricultural industries, electricity generation and construction[2].

The novelty in this dissertation is the updated data ranging from 2001–2019, allowing me to estimate most of phase III. Many studies rely on estimated emissions calculated from, for instance, energy consumption from various energy carriers (Petrick and Wagner 2014; Löschel, Lutz, and Managi 2019). This adds a certain level of uncertainty. Two honorary exceptions are Bel and Joseph (2015) and Klemetsen, Rosendahl, and Jakobsen (2020) in using verified emissions reported by each firm. This paper takes the same approach by using emissions data reported by installations.

The main drawback of this dataset is the lack of plant-level economic or energy variables. This makes matching on common characteristics, particularly the selection criteria for the EU ETS, more difficult. A more detailed discussion on how this is dealt with is covered in chapter 5.

1. The data in part consists of data made available for this project by NEA, and is in part downloaded from the Norwegian PRTR website. on 3 July 2020.
2. See Statistics Norway for a description of SIC 2007 for Norwegian firms.

Figure 4.1: Number of ETS and non-ETS firms

Note: The figure plots the number of ETS and non-ETS firms in the dataset in each year.

Firms

In order to internalise some spillovers within firms, I aggregate up from the plant level to the firm level. Firms (as opposed to companies) is the lowest level of legal unit in the Norwegian firm registry. The aggregation therefore entails little loss of detail, but facilitates interpretation. The unbalanced dataset covers 345 firms, 19 years and 3647 firm-year observations. The firms belong to 106 different five-digit industries and 23 different two-digit industries. As shown in figure 4.1, there are 51 ETS firms in phase II and 69 ETS firms in phase III.

4.2 Variables

Emissions

I observe a yearly measure of greenhouse gas (GHG) emissions for each firm, as tonnes of CO_2 equivalents (CO_2eq). This includes all emissions covered by the Kyoto protocol, including CO_2 from fossil fuels (but not biofuels) [3]. I also observe tonnes of carbondioxide (CO_2) emissions, CO_2 from fossil fuels and

3. In particular, this measure also includes CH_4, CF_4, C_2F_6, SF_6, N_2O and HFCs (hydrofluorocarbons).

nitrous oxide (N_2O).

	mean	sd	count
GHG	65535.91	206231.4	3647
CO2	63900.16	182855.2	3556
N2O	29.01894	268.913	1992
CO2 from fossil fuels	59681.4	179760.9	3483

Table 4.1: Descriptive statistics for emissions variables

Note: This summarises the mean, standard deviations and number of firm-year observations of the emissions variables in my dataset (each measured in tonnes), for both ETS and non-ETS firms from 2001–2019.

Table 4.1 shows that the majority of GHG emissions in my sample come from CO_2 , mainly from burning of fossil fuels. However, N_2O is about 300 times as powerful a greenhouse gas as CO_2 and is a considerable share of the emissions prior to 2008, but become less important in later years. Perfluorocarbons (PFCs), while also covered by the EU ETS, are not observed in my dataset. They can, however, be considered one of the key sources of the remaining variation in GHG after controlling for CO_2 and N_2O .

ETS firms have on average much higher emissions than non-ETS firms, see figure 4.2. Phase II ETS firms have four times higher GHG emissions than Phase II non-ETS firms. Meanwhile, Phase III ETS firms have nearly 40 times higher GHG emissions than Phase III non-ETS firms. This shows that the inclusion of aluminium firms in phase III causes the difference between ETS and non-ETS firms to explode, resulting in large differences between the firms that are and are not covered by the ETS from 2013 onwards.

There was a decline in emissions from 2005 to 2007 due to technological changes curbing N_2O emissions from the manufacturing industry (NEA 2010). This is particularly apparent for the ETS-firms in figure 4.2, but the changes impacted non-ETS firms too. There also seems to be some drop in mean emissions, particularly for phase III-firms, around the time of the recession[4]. After 2009, mean emissions increase somewhat, but never reach pre-2006 lev-

4. Some of the impact of the recession is likely to be reflected in firms dropping out of the market, not only through reduced output and hence emissions from surviving firms.

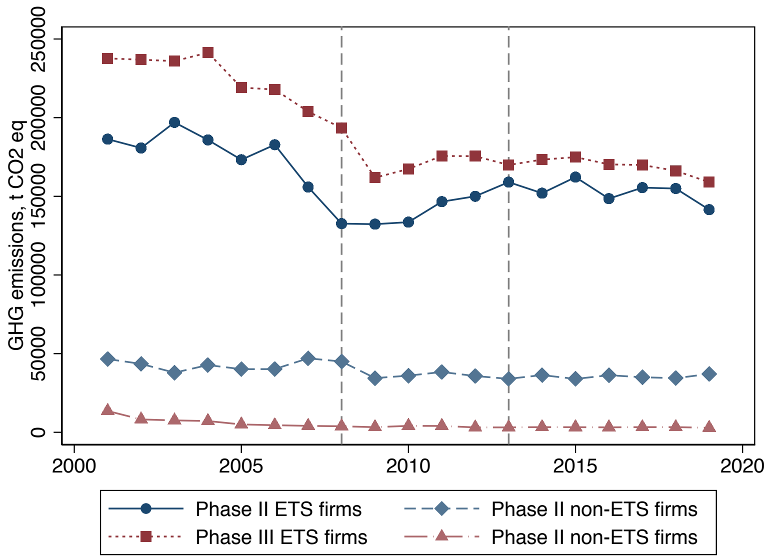

Figure 4.2: Mean emissions of CO2 equivalents

Note: The plot shows the different yearly means of emissions of GHG for ETS and non-ETS firms in phases II and III, respectively. Phase III-firms include firms that were also covered in phase II, while non-ETS firms in phase II also include the firms that became covered by the ETS from phase III.

els. There is a similar trend for firms that were not covered in phase II, while firms that are not covered in phase III saw a marked decline in GHG emissions from 2001 to around 2009. For all four categories, mean emissions stabilised from around 2011.

As a proxy of the production capacity of firms prior to introduction the EU ETS I apply firm emissions from 2001 and 2002[5]. Treatment in the EU ETS depends on the industry the plant belongs to and certain industry-specific thresholds of either thermal input, production output or capacity. While I observe the specific industry code, I do not observe the output of the firm or plant. This induces an omitted variable bias in the regression, as output and thermal input clearly impact emissions as well as regulated status. To abate this problem I propose the use of lagged emissions as 'proxy' variables for the ETS regulation criteria. Wooldridge (2010) establishes two assumptions that must hold for a valid proxy variable, (4.1) and (4.2):

$$\mathbb{E}(y|\boldsymbol{x}, \boldsymbol{q}, \boldsymbol{z}) = \mathbb{E}(y|\boldsymbol{x}, \boldsymbol{q}) \tag{4.1}$$

5. See Holzer et al. (1993) for an example of a paper that applies the logged dependent variable as a proxy.

where $\boldsymbol{x}$ denotes the vector of explanatory variables, $\boldsymbol{q}$ the vector of omitted variables and $\boldsymbol{z}$ the vector of proxy variables. (4.1) means that, given the explanatory and omitted variables, the conditional mean of emissions should not depend on the value of the proxy variable. Wooldridge (2010) calls (4.1) the redundancy criterion. The second criterion

$$\begin{aligned} \boldsymbol{q} &= \theta_0 + \boldsymbol{z}'\theta_1 + r \\ \mathrm{Cov}(x_j, r) &= 0 \qquad j = 1, 2, ..., K \end{aligned} \tag{4.2}$$

meaning that $\boldsymbol{z}$ is sufficiently strong in explaining $\boldsymbol{q}$ to ensure that the explanatory variables are no longer partially correlated with the omitted variable, once the proxy has been included. In the remaining analysis I assume that (4.1) and (4.2) hold.

Industry affiliation

I control for industry-specific regulation and other characteristics of the industries by including industry dummies. To retain sufficient variation I apply two-digit industry level codes.

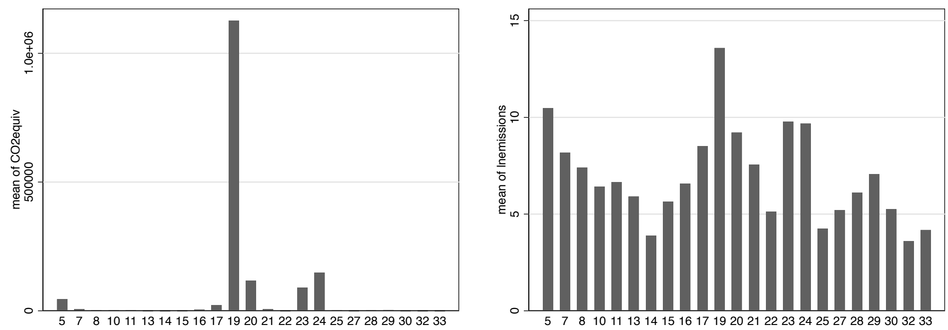

(a) Mean GHG emissions, by industry code

(b) Mean of logged GHG emissions, by industry code

Figure 4.3: Mean emissions by industry code

Note: Panel (a) illustrates how mean emissions vary a great deal across industry codes. The means of the natural logarithm of firm emissions, plotted in panel (b), are not as dominated by outliers.

Figure 4.3 plots mean GHG emissions by industry code. Industry 19, manufacturing of refined petroleum products, stands out as the industry with the largest mean emissions in panel A. Since 2008, all these firms have been regulated. In my analysis I apply logged GHG emissions as my dependent variable in order to approximate a normal distribution. Panel B illustrates how logged GHG emissions are more comparable across industries.

5 Methodology

Identifying the impact of the EU ETS on emissions requires disentangling the impact of the EU ETS from variables such as macroeconomic conditions (e.g. business cycles) and wider technology development. Typically, causal EU ETS studies exploit the variation caused by some firms being treated and others remaining untreated (Martin, Muûls, and Wagner 2016). This paper draws inspiration from the large majority of previous EU ETS research which applies difference-in-differences and linear regressions, often with various types of matching (see table 3.1). In this chapter, I assess the implications and credibility of the assumptions the methods rest on. I account for the non-randomness of the selection criteria (industry affiliation and production or input capacity) through propensity score matching.

As a non-member state of the EU, the most relevant counterfactual to Norwegian ETS firms is Norwegian non-ETS firms participating in a European product and energy market where the EU ETS still operates. I assume that ETS treated firms would have faced regulation similar to that of non-treated firms in the absence of ETS regulation, see section 2.3. In doing so, I derive estimators that should be interpreted as the impact of the EU ETS on Norwegian treated firms, compared to the counterfactual of the regulatory context of non-treated firms. This gives the estimates a somewhat different interpretation to that of Klemetsen, Rosendahl, and Jakobsen (2020), which I return to in chapter 6.

5.1 Treatment effect

When setting out to establish a causal effect, the ideal research design is a randomised control trial. In most economics research, this is infeasible and unethical. Economists therefore resort to natural 'as-if random' experiments and attempt to correct for selection bias and non-randomness in assignment of treatment. I only observe the realised outcomes – the difference in means between the treated and untreated firms – not the counterfactuals. Following the tradition of the Rubin causal model and Angrist and Pischke (2009), I denote the potential outcome of individual i as Y_{1i} if individual i is treated, and Y_{0i} if i is untreated. Whether or not the firm belongs to the treatment group is a dummy I denote by ETS_i. This distinction is needed because we differentiate between the observed outcomes $\mathbb{E}(Y_{1i}|ETS_i = 1)$ for treated firms and $\mathbb{E}(Y_{0i}|ETS_i = 0)$ for untreated firms, and the unobserved counterfactuals $\mathbb{E}(Y_{0i}|ETS_i = 1)$ and $\mathbb{E}(Y_{1i}|ETS_i = 0)$. The average treatment effect on the treated (ATT) is then given by

$$\tau_{ATT} = \mathbb{E}[Y_{1i} - Y_{0i}|ETS_i = 1] \tag{5.1}$$

(5.1) illustrates that I would ideally like to estimate is the difference between what happens to the individual that is treated, and the counterfactual effect if they were not treated. While the average treatment effect (ATE) could be estimated instead, this requires stronger assumptions that cannot be ensured to hold in this sample (Wooldridge 2010). The limitation to ATT allows estimation based on weaker and more reasonable assumptions, without much loss of generality.

A natural starting point is the observed difference in means between the treatment and control group. If there are differences (observed or unobserved) between the firms that receive and do not receive treatment, there is a selection bias (Angrist and Pischke 2009):

$$\begin{aligned} \underbrace{\mathbb{E}[Y_{1i}|ETS_i = 1] - \mathbb{E}[Y_{0i}|ETS_i = 0]}_{\text{observed difference in means}} = &\underbrace{\mathbb{E}[Y_{1i}|ETS_i = 1] - \mathbb{E}[Y_{0i}|ETS_i = 1]}_{\tau_{ATT}} \\ &+ \underbrace{\mathbb{E}[Y_{0i}|ETS_i = 1] - \mathbb{E}[Y_{0i}|ETS_i = 0]}_{\text{selection bias}} \end{aligned} \tag{5.2}$$

This selection problem is a classic one in econometrics, and no less so in empirical environmental economics. It is clear from (5.2) that the selection problem only arises when there is a correlation between the selection criteria for treatment and the outcome variable. This goes to the core of the identification problem of ETS studies: ETS regulation is determined by the industry affiliation and production capacity of the firm, while emissions largely depend on technology and production volumes. There is therefore reason to expect a relatively large selection bias in the absence of a sophisticated identification strategy. In order to develop a credible identification strategy, I present the three main assumptions highlighted in Wooldridge (2010) for estimation of ATT:

The potential (counterfactual or realised) outcomes for each individual i are independent of whether or not another individual j is treated (Angrist, Imbens, and Rubin 1996). (5.3)

(5.3) is commonly known as the Stable Unit Treatment Value Assumption (SUTVA). SUTVA is a key assumption in requiring that treatment affects only the firms that receive the treatment, such that untreated firms remain unaffected. In the ETS case, this means that there cannot be any spillovers to non-ETS firms. However, the ETS has led to increased energy prices in much of the EU, impacting the trade-off between clean and dirty inputs for non-ETS and ETS firms alike (see e.g. Hintermann, Peterson, and Rickels 2016). Further, increased innovation caused by emissions trading has often been seen as a key explanatory force for the seemingly improved economic performance of firms (see e.g. Klemetsen, Rosendahl, and Jakobsen 2020). If this makes affordable technology more easily available to non-ETS firms, this can lead to spillovers reducing emissions of both ETS and non-ETS firms. It is therefore not possible to ensure that SUTVA holds with certainty in this analysis. However, the treatment is applied strictly to the ETS firms only and would not cause any direct impact on non-ETS firms.

In order to solve the selection problem, researchers often require that the necessary assumptions hold in the mean conditional on observables, such as the selection criteria. This property is captured by the ignorability assumption:

$$\mathbb{E}(Y_{0i}|\boldsymbol{x_i}, ETS_i) = \mathbb{E}(Y_{0i}|\boldsymbol{x_i}) \tag{5.4}$$

This assumption requires that, conditional on the covariates $\boldsymbol{x}$, treatment is independent of the outcome variable in its mean (Wooldridge 2010). It is commonly called 'selection on observables', and requires that the only unobservables treatment ETS_i is allowed to depend on are uncorrelated with treatment y.

The final assumption is the overlap assumption,

$$P(ETS_i = 1|\boldsymbol{x_i}) \in [0, 1) \quad \forall \boldsymbol{x_i} \in \boldsymbol{\mathcal{X}} \tag{5.5}$$

where $\boldsymbol{\mathcal{X}}$ denotes the support of $\boldsymbol{x}$. The overlap assumption for ATT requires that all observations included in the estimation have a probability of receiving treatment that is less than one (conditional on the covariates). The main problem with the lack of overlap is that the results of estimation, particularly τ_{ATT}, is not identified at the values of $\boldsymbol{x}$ for which treatment is perfectly predicted. This means that τ_{ATT} cannot be generalised to hold for the values of $\boldsymbol{x}$ for which there is insufficient overlap (Wooldridge 2010).

5.2 Identification

In an attempt to identify the causal effect satisfying assumptions (5.3), (5.4) and (5.5), I apply several identification strategies. At the core of my identification strategy lies propensity score matching, which is then applied in 1) a difference-in-differences framework and 2) a linear panel data regression. I run a set of specifications on three samples – the entire sample, the sample on the common support and the inverse propensity score weighted sample on the common support. In using multiple strategies I aim to assess whether the results are consistent across different identification strategies.

In applying these, I also disregard some other candidate strategies. For instance, regression discontinuity design (RDD) might seem a reasonable approach given the structure of the EU ETS selection criteria. However, both Klemetsen, Rosendahl, and Jakobsen (2020) and Petrick and Wagner (2014) consider RDD infeasible. In my case, RDD would serve to estimate the local effect of firms very close to the regulation threshold consistently, but this would both render the sample very small, as well as have little to no generalisability beyond this local region. Further, RDD would identify an effect different to that of the other strategies, and would therefore not be suitable for comparisons with previous research.

I conduct difference-in-differences (DiD) and linear regression estimation. In addition, I apply what Angrist and Pischke (2009) define as regression adjustment, where I trim the sample to the firms that are on the common support and condition on the covariates used to predict treatment status. This ensures that both (5.5) holds, and, if the regression specification is correct, that (5.4) holds. Further, I apply regression adjustment with propensity score weighting. This is doubly robust in the sense that it suffices that either the regression model or the propensity score model is correctly specified (Angrist and Pischke 2009), making it more likely that (5.4) is satisfied. It also constitutes a very simple test of the regression specification – if results deviate strongly, it may suggest that the regression model is misspecified.

Macroeconomic trends and business cycles, e.g. the 2008/2009 recession, are likely to explain some of the variation in emissions for both ETS and non-ETS firms. In 2014, the oil price plummeted, causing a decline in oil and gas extraction and in demand for the industries supplying technology, services and products for the oil industry. Unemployment saw a sharp increase, particularly in Western Norway. While business cycles impact all firms, they do so differently. This justifies the application of a difference-in-differences model in order to allow for group-level fixed effects (Angrist and Pischke 2009). In doing so, I assume that the parallel trend assumption holds, such that logged emissions of the control and treatment groups would have followed the same trend in the absence of regulation (conditional on the covariates). This also requires that treated and untreated firms are impacted the same by these macroeconomic shocks.

As the counterfactual is not observed this cannot be tested directly. However, I can investigate the parallel trend assumption by visual inspection of plots of logged emissions for the three different specifications (entire sample, common support and inverse propensity weighted sample). I only observe the firms from 2001 onwards, and therefore do not observe too many years prior to announcement of the EU ETS. Any deviation from the parallel trend assumption therefore cannot be taken as a clear violation from the parallel trend assumption, but would suggest that the results should be interpreted with caution.

Figure 5.1 plots the mean of log GHG emissions, by treatment status and year. In the full sample, panel (a), the trend is the same as in chapter 4: ETS firms have far higher mean emissions than non-ETS firms, particularly in phase III. From 2002 to 2004, there is a decline in logged emissions from non-ETS firms, but not from ETS firms. After 2005 (when phase I begins), non-ETS firms see a slight increase in emissions, while ETS firms do not. There is a slight increase in mean non-ETS logged emissions from 2010 to 2011, but otherwise, the firms seem to follow a relatively similar trend even

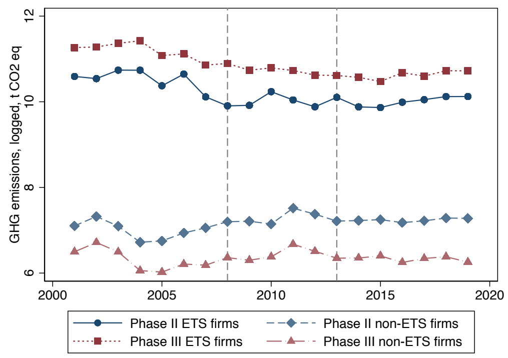

(a) Full sample

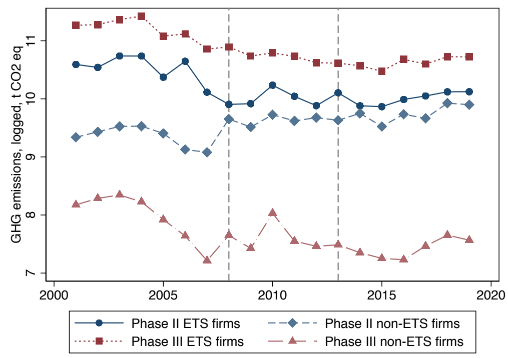

(b) Common support

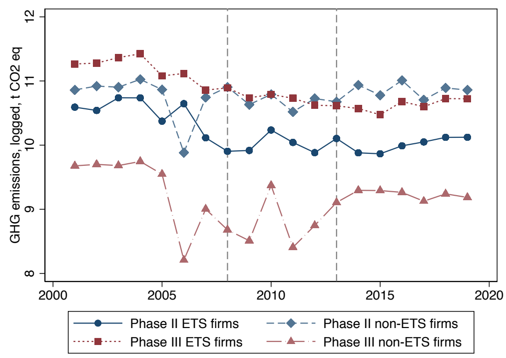

(c) Inverse propensity weighted sample

Figure 5.1: Mean of logged GHG emissions by ETS treatment status

Note: The figure plots the mean of the yearly natural logarithm of GHG emissions of firms by their ETS treatment status in phase II and phase III. Panel (a) plots mean log emissions for all 345 firms in the dataset. Panel (b) plots all 119 firms on the common support. Panel (c) plots the inverse propensity weighted mean of logged emissions for the 119 firms on the common support.

after treatment is introduced in 2008 and 2012, respectively. For the firms on the common support, panel (b), the means seem more volatile, likely in part because there are fewer firms in this sample. The log means are now more similar for ETS and non-ETS firms, both in phase II and III. This suggests that the matching variables described in chapter 5 are good predictors of logged emissions in future years. The trends prior to 2005 are similar, lending support to the parallel trend assumption. Firms that are non-ETS in phase III see a decline in logged emissions not shared by ETS firms. Inverse propensity weighting (panel (c)) reduces the absolute difference in log means, particularly in 2001 and 2002. There now seems to be a sharper increase in log means for ETS firms than for non-ETS firms prior to 2005. In the following few years, there is a sharp decline for non-ETS firms not reflected in the trends of the

ETS firms. After 2008, ETS and non-ETS firms seem to follow a very similar trend.

Propensity scores

The EU ETS presents a textbook example of selection bias: Only units with a certain minimum capacity and belonging to specific industries are treated. In order to ensure that the ignorability assumption (5.4) holds, I must control for the difference in probability of recieving treatment caused by these selection criteria. Otherwise, my estimator for τ_{ATT} would capture the impacts of production capacity and industry technology as well as the impact of the EU ETS.

Treatment status of Norwegian firms in the EU ETS is determined by the industry and production capacity of each firm in the period 1998–2001. In the treatment literature, the probability of treatment conditional on observed covariates is typically called the propensity score and is given by (5.6) (Angrist and Pischke 2009):

$$p(\boldsymbol{x}) = P(ETS = 1|\boldsymbol{x}) \tag{5.6}$$

As is conventional, the propensity score of firm i is estimated using a logit model and the Stata programme *psmatch2* by Leuven and Sianesi (2003). Perhaps surprisingly, taking into account the uncertainty in estimating the propensity score can lead to lower estimates for the regression standard error (Wooldridge 2010). The disadvantage of the *psmatch2* programme is therefore that it does not allow me to exploit the fact that the propensity score is estimated in my further analysis, despite this often being more efficient. This choice of programme is perhaps at the cost of efficiency, but not unbiasedness. While bootstrapping can be applied to limit the standard error, it is unclear in the literature whether this is valid and I therefore avoid it.

Following Klemetsen, Rosendahl, and Jakobsen (2020), I match observations with up to ten of its closest neighbours, with three neighbours as a robustness test. I modify the assumptions to be conditional on the propensity score (i.e. a function of the covariates), rather than the covariates themselves. Therefore, in addition to SUTVA, the assumptions are given by

$$\text{Ignorability:} \qquad \mathbb{E}(Y_{0i}|p(\boldsymbol{x_i}), ETS_i) = \mathbb{E}(Y_{0i}|p(\boldsymbol{x_i})) \tag{5.7}$$

$$\text{Overlap:} \qquad P(ETS = 1|p(\boldsymbol{x})) \in [0, 1) \quad \forall \boldsymbol{x_i} \in \boldsymbol{\mathcal{X}} \tag{5.8}$$

Angrist and Pischke (2009) discuss the following paradox of propensity score methods: Estimators based on the propensity score are asymptotically less efficient than estimates based on the covariates themselves. However, propensity score methods have good finite-sample properties (Angrist and Hahn 2004). By restricting the propensity score through nonparametric methods using prior knowledge of the criteria for treatment, the researcher is able to implement additional information and experience efficiency gains, leading to improved finite-sample results (Angrist and Pischke 2009).

In order to identify the average treatment effect on the treated, ATT, I conduct exact matching within the two-digit industry code to ensure that all neighbours are within the same industry, and match on lagged log GHG and CO_2 emissions (2001 and 2002). I use a linear form of emissions because treatment is determined by whether or not capacity exceeds a certain threshold. By conducting exact matching I only consider variations in emissions within industries, which is appropriate when thresholds are industry-specific.

For treated firms, the propensity score weight w is 1, because I estimate the average treatment effect on treated firms (ATT),

$$w(\boldsymbol{x}|ETS = 1) = 1 \tag{5.9}$$

Untreated firms are weighted by the estimated likelihood of treatment,

$$w(\boldsymbol{x}|ETS = 0) = \frac{\hat{p}(\boldsymbol{x})}{1 - \hat{p}(\boldsymbol{x})} \tag{5.10}$$

In other words, the analysis puts more weight on non-treated firms with higher probability of treatment, and compares them with the treated firms. This obtains the inverse propensity weighted ATT, which provides consistent estimates for τ_{ATT} under assumptions (5.7) and (5.8) (Angrist and Pischke 2009).

6 Results

In order to compare and contrast the findings of different models, I implement models with a gradually increasing number of controls. First, I run linear and difference-in-differences (DiD) regressions on the entire, unweighted sample. I expect the models to respond noticeably to conditioning on the selection criteria and industry- and time-fixed effects. This is also why I expect to see differences in the results from the linear regression and difference-in-differences models, as DiD controls for time trends and group-fixed effects as well as those specified as control variables[1]. A key reason for applying DiD in this analysis is to isolate from the effects of demand-side variation in the economy, e.g. increased demand for industry goods due to macroeconomic growth or shocks to the economy. I do not observe a measure of production value for each firm, nor would I be able to condition on this, as changes in production volumes is one of the channels through which firms can adjust their emissions in response to regulation. If the firms are sufficiently similar and can be expected to follow the same trends, these effects are captured instead by the DiD specification.

Second, I conduct propensity score matching and run the same models on a sample trimmed to those firms that are on the common support $\boldsymbol{\mathcal{X}}$. As there are major differences between the ETS and non-ETS samples, I expect to see some changes in the results. Some of the improvement from this approach can also be achieved by conditioning on certain variables. For instance, when I control for industry affiliation I only do the analysis on industries for which there are both ETS and non-ETS firms in the same period. However, by limiting the analysis to the firms on the common support, I do not consider firms that receive treatment with certainty. This means that I am left with 1401 firm-year observations. This is also the main drawback of limiting analysis to the common support – only a few industries are included in this

1. Klemetsen, Rosendahl, and Jakobsen (2020) and others conduct their linear regression analysis using firm-fixed effects. In this model, analysis with firm-fixed effects causes too much collinearity and is infeasible.

sample, and the findings cannot be generalised beyond the common support. It also comes at the cost of precision, in the sense that the sample is reduced. The choice of applying analysis to the common support is therefore a tradeoff between generalisability and precision, and unbiasedness and consistency of results.

Finally, I weight the sample on the common support by the inverse propensity score as described in chapter 5. The implication is that similar firms receive more weight than dissimilar firms. The DiD approach rests on the parallel trend assumption, and IPSW makes this assumption more credible.

Following the reasoning in Abadie et al. (2017) I cluster standard errors on a five-digit industry level because treatment is largely dependent on the industry affiliation of firms. In an alternative specification I run the same regressions with firm-level clustering, which returns smaller standard errors. Before trimming the sample, I have a large number of five-digit industries (106), but this goes down to 29 after trimming the sample. Clustering is therefore to some extent at the cost of efficiency within the trimmed sample (Abadie et al. 2017).

6.1 Difference-in-differences

The population regression DiD model is given by (6.1):

$$\ln E_{it} = \beta_0 + ETS_{ip}\beta_1 + \alpha_j + \eta_t + \alpha_j' x_{i,pre} + \gamma_e\beta_2 + \varepsilon_{it} \tag{6.1}$$

The dependent variable, $\ln E_{it}$, is the natural logarithm of the GHG emissions of firm i in year t. ETS_{ip} is a 1×2 dummy vector for whether firm i is covered by the ETS in phase p. The 2×1 vector of coefficients β_1 is therefore the coefficient on the impact occurring in phase II and III. α_j is a dummy for two-digit industry affiliation, and allows me to control for industry-fixed effects. $x_{i,pre}$ is a vector of logged firm emissions of GHG and CO_2 in 2001 and 2002, used as proxy variables for production or input capacity prior to the announcement of Norwegian participation in EU ETS. The interaction term with the industry dummies thus captures the selection criteria. η_t captures year-fixed effects such as overall economic growth, while γ_e is a dummy for whether or not the firm belongs to the treated or untreated group. The inclusion of γ_e is what differs between the DiD and the linear regression model in (6.2). ε_{it} is the error term, which is assumed to be uncorrelated with

ETS_{ip}. The model is specified twice, with and without the industry-capacity interaction term, to investigate the sensitivity of the model to the selection criteria. The findings are reported in table 6.1.

The model in table 6.1 implements (6.1) with and without the interaction term I apply for the selection criteria on three different samples. Columns (3) and (4) are the trimmed sample on the common support, $\boldsymbol{\mathcal{X}}$. Column (5) weights the control group by the inverse propensity score, and column (6) is 'doubly robust' – it weights the control group by the inverse propensity score in addition to controlling for the selection criteria.

Overall, the estimated impact occurring in phase II is negative, suggesting that the ETS caused phase II firms to have around 20 % (or more) lower GHG emissions in phase II than non-ETS firms. Further, the estimated impact occurring in phase III is positive, suggesting that phase III-firms had around 30–50 % higher GHG emissions in phase III than they would have had in the absence of the ETS. The estimates for phase II are significant at least at the 10 % level in all columns except (6). Notably, the estimates become lower and less significant with more conservative specifications, but remain negative. The estimates for phase III are significant in all columns except (4), and seem to be of similar magnitude across specifications. The estimates for the impact of the EU ETS are generally higher in specifications (1), (3) and (5), which are likely inflated by capturing some of the impact of the selection criteria, due to selection bias.

If the interaction term $\alpha_j' x_{i,pre}$ absorbed the entire time-invariant correlation between GHG emissions and the regulated status, I would have expected to see a near-zero coefficient on γ in columns (4) and (6). While the coefficients are much lower than in the absence of the selection criteria, it is still significant at the 1 % and 5 % level, respectively. The magnitude of the estimated impacts of the ETS occurring in phase II and phase III are also notably smaller, suggesting that there is a group-fixed effect of being an ETS firm that is not in its entirety captured by the selection criteria term. However, the large drop in the coefficients on γ when including $\alpha_j' x_{i,pre}$ suggests that the interaction between industry affiliation and 2001 and 2002 emissions to a large extent captures the time-invariant impact of being an ETS-treated firm.

The consistently negative estimates in phase II lend support to the findings in Klemetsen, Rosendahl, and Jakobsen (2020), who found around 30

% abatement for ETS-firms in phase II. Further, the consistently positive estimates in phase III suggest a positive impact of ETS regulation in phase III, compared to the counterfactual scenario. Implications of this are further discussed in chapter 7.

6.2 Linear regression model

Linear panel data regression is adopted as an alternative specification to the DiD model. The population regression model is similar to that in (6.1), but does not contain the ETS group dummy γ_e (the remainder of the variables and coefficients are explained at the beginning of section 6.1):

$$\ln E_{it} = \beta_0 + ETS_{ip}\beta_1 + \alpha_j + \eta_t + \alpha_j' x_{i,pre} + \zeta_{it} \tag{6.2}$$

The model is, as such, less conservative than the DiD model. Instead of resting on the parallel trend assumption it assumes orthogonality of the group fixed effects and treatment. The results are given in table 6.2.

Columns (1), (3) and (5) run the analysis without controlling for the selection criteria or a group dummy variable. These columns have estimates that seem very inflated compared to those in table 6.1, with an estimated impact in phase III of between 170 % and 310 %. The majority of this difference between the linear regression and DiD models disappears once the selection criteria are added. This suggests that the phase III coefficient captures some of the time-invariant correlation with selection into the treatment group. Further, the results in table 6.1 suggests that there is indeed a group-fixed effect not entirely captured by the selection criteria. The linear regression model (6.2) might therefore be less suitable for explaining the impact of the EU ETS than the DiD model (6.1). The linear regression model is prone to omitted variable bias and its results should be interpreted with caution.

In sum, however, all twelve specifications return the same sign on the estimates for the two phases. While the findings in table 6.1 rest very heavily on the parallel trend assumption, the results in table 6.2 all indicate a reduction in emissions for ETS-firms in phase II and an increase in phase III, relative to the counterfactual. The estimates in models (3)–(6) suggest that at least in the four two-digit industries on the common support – manufacturing of paper and pulp, chemicals, non-metallic mineral products and basic metals – Norwegian ETS firms saw a decline in emissions relative to Norwegian non-

ETS firms in phase II, and an increase in phase III. Note that this is heavily dependent on the assumptions set out in chapter 5 being satisfied, and its implications are further analysed in the following chapter.

	(1) lnemissions	(2) lnemissions	(3) lnemissions	(4) lnemissions	(5) lnemissions	(6) lnemissions
EU ETS in phase II	-1.241***	-0.532**	-1.148***	-0.410**	-0.674*	-0.221
	(0.420)	(0.215)	(0.415)	(0.187)	(0.334)	(0.219)
EU ETS in phase III	0.496**	0.527**	0.593***	0.307	0.563**	0.397*
	(0.236)	(0.226)	(0.198)	(0.182)	(0.210)	(0.205)
ETS firm (γ)	4.189***	0.927***	4.066***	0.760***	1.937***	0.569**
	(0.414)	(0.238)	(0.492)	(0.194)	(0.308)	(0.272)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Selection criteria	No	Yes	No	Yes	No	Yes
Trimmed sample	No	No	Yes	Yes	Yes	Yes
Weighted	No	No	No	No	Yes	Yes
Observations	3539	2352	1401	1401	1401	1401

Standard errors in parentheses

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

Table 6.1: Difference-in-differences

Note: The table shows the difference-in-differences estimated model with and without controls for the selection criteria (an interaction term between a two-digit industry dummy and logged firm emissions of CO_2 and greenhouse gases in 2001 and 2002). The dependent variable is the natural logarithm of greenhouse gas emissions. The dummy variables for phase II and phase III regulation capture whether or not the firm was ETS-regulated in the given phase, and thus the impact of the ETS occurring in that phase. The group dummy variable γ captures whether the firm belongs to the ETS group or not. The trimmed sample is the result of a propensity score matching with up to 10 neighbours on the selection criteria, where only those firms on the common support are kept. The weighted sample weights the observations by the inverse propensity score resulting from the same matching. The data are collected from the Norwegian Environmental Agency (NEA) and cover 345 firms from 2001–2019.

	(1)	(2)	(3)	(4)	(5)	(6)
	lnemissions	lnemissions	lnemissions	lnemissions	lnemissions	lnemissions
EU ETS in phase II	-0.675	-0.398*	-1.037*	-0.326*	-0.504	-0.109
	(0.526)	(0.202)	(0.563)	(0.175)	(0.357)	(0.168)
EU ETS in phase III	3.143***	0.810***	3.536***	0.602**	1.725***	0.647**
	(0.591)	(0.258)	(0.614)	(0.227)	(0.423)	(0.308)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Selection criteria	No	Yes	No	Yes	No	Yes
Trimmed sample	No	No	Yes	Yes	Yes	Yes
Weighted	No	No	No	No	Yes	Yes
Observations	3539	2352	1401	1401	1401	1401

Standard errors in parentheses

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

Table 6.2: Linear regression

Note: The table shows the linear regression panel data estimated model with and without controls for the selection criteria (an interaction term between a two-digit industry dummy and logged firm emissions of CO_2 and greenhouse gases in 2001 and 2002). The dependent variable is the natural logarithm of greenhouse gas emissions. The dummy variables for phase II and phase III regulation capture whether or not the firm was ETS-regulated in the given phase, and thus the impact of the ETS occurring in that phase. The group dummy variable γ captures whether the firm belongs to the ETS regulated group or not. The trimmed sample is the result of a propensity score matching with up to 10 neighbours on the selection criteria, where only those firms on the common support are kept. The weighted sample weights the observations by the inverse propensity score resulting from the same matching. The data are collected from the Norwegian Environmental Agency (NEA) and cover 345 firms from 2001–2019.

7 Discussion

Chapter 6 presented the results from the analysis of the data. The estimates all point in the direction of reduced GHG emissions for ETS firms in phase II, and increased GHG emissions in phase III, compared to the non-ETS firms. While the methodological caveats place some limitations on interpretation of the coefficients, the results of the analysis presented in chapter 6 indicate a relatively clear direction. In phase II, the results are similar to those of Klemetsen, Rosendahl, and Jakobsen (2020) and Wagner et al. (2014). In phase III, there are consistently positive and relatively high estimates, even in the most robust specifications. While this may be attributable to research design and data limitations, it also suggests that there might have been more abatement within industrial firms in Norway without EU ETS regulation. This cannot immediately be taken to mean that industrial emissions in Europe would have been lower, however, as Norwegian firms have participated to increasing the demand for EUAs, thereby bringing up the price. In this chapter I discuss various possible explanations for both of these findings.

The findings in this dissertation lend additional evidence to the growing literature finding negative impacts of phase II regulation in several countries, see chapter 3. To explain their findings in phase II, Klemetsen, Rosendahl, and Jakobsen (2020) focus primarily on the price of EUAs. In the early years of phase II, the prices of EUAs were at an historically high level, see figure 2.2. Low EUA prices likely created stronger incentives for abatement. However, as an increasingly large surplus of EUAs started flooding the market at the same time as the recession hit the European economies, prices of EUAs dropped and did not start increasing again until 2018. The option to bank allowances enables firms to optimise dynamically such that they stock up on allowances for the future if they expect the price to rise. Therefore, high EUA prices prior to the recession can be interpreted as firms expecting equally high prices in the future, too. This may have created incentives to invest in innovation in cleaner technologies, causing the abatement to continue further into phase II.

The key novelty of this paper is extending the analysis to phase III, where I find consistently positive estimates for the impact of the EU ETS on Norwegian ETS firms compared to non-ETS firms. There can be multiple reasons for these findings. The estimates in chapter 6 only extend to the industries on the common support, and may not be generalisable beyond that. To the extent that the results hold and can be generalised, I discuss along two dimensions: 1) Whether the low EUA prices in phase III suggest that the EU ETS has been inefficient in delivering abatement in phase III, and 2) Whether there are particular characteristics of Norwegian manufacturing and mining firms that can explain these findings.

A major caveat of all methods applied in this dissertation is the lack of generalisability of the ATT beyond the industry codes on the common support. As Wooldridge (2010) highlights, ATT is only identified where there are firms with a probability of treatment strictly below 1. The combination of relatively few industrial firms in Norway and the thresholds used to determine the regulated status of firms within the ETS mean that in the trimmed sample, only four two-digit industries are represented: manufacturing of paper, chemicals, non-metallic mineral products (e.g. glass, cement) and basic metals (e.g. steel). These make up a total of 29 different five-digit industries. The results therefore cannot be generalised to hold for mining and quarrying or other manufacturing industries – despite these firms being part of the full sample.

The consistently positive estimates on the impact of the ETS in phase III raise the question of whether the EU ETS has efficiently caused abatement at all, at least in the industries listed above. Phase III is characterised by more harmonised rules across EU ETS member states and limited access to linked allowance markets. Meanwhile, there also was a low price on EUAs for the majority of the phase, see figure 2.2. A key assumption for MAC to impact the EUA price is that the cap is 'tight', i.e. exceeds the business-as-usual emissions (see chapter 2). The recognised surplus of allowances in much of phase III suggests that the cap might in fact have been 'slack'. Little incentive will have been provided for Norwegian ETS firms to reduce their emissions due to the low allowance price. For the majority of phase III the price of EUAs therefore provided little incentives for abatement. Further, this may have caused expectations of low future EUA prices, creating few incentives for innovation and implementation of green technologies, limiting both the long-and short-term impact of the ETS on abatement.

On the other hand, the positive estimates can be explained by abatement occurring elsewhere in the ETS due to differences in MAC. The majority of phase III firms were also regulated in phase II. However, the Norwegian phase III-firms that were not regulated in phase II largely belong to aluminium production and received incentives before entering the ETS (Klemetsen, Rosendahl, and Jakobsen 2020), see section 2.3. If these measures had any impact, Norwegian aluminium industry may have already harvested much of the low-hanging fruit before entering phase III. Further, there has been an overlap between the ETS and climate policies in other member states and the EU, such as incentives for renewable energy generation, which may have effectively brought down the MAC of affected firms (Ellerman, Marcantonini, and Zaklan 2016). If MACs were significantly lower elsewhere in the EU ETS, this will have brought down the price of EUAs despite the cap being 'tight'. Norwegian firms will then have had incentives to pay for abatement to occur elsewhere in phase III, bidding up the price of EUAs. While this does not speak to the EU ETS lacking capability to deliver emissions abatement, it suggests that Norwegian firms have made use of the flexibility mechanisms of the EU ETS rather than bearing their own abatement costs.

While low EUA prices may have provided little incentive for abatement, it cannot explain the *positive* coefficients estimated. However, abatement incentives may have been stronger for non-ETS firms. For instance, the CO_2 tax has increased several times during phase III, and multiple climate policies aimed at non-ETS firms have been in place (NEA 2010, 2020). As non-ETS firms have not had the option of buying allowances from firms with lower MAC, they may have had stronger incentives to conduct abatement. Assuming a 'slack' cap and stronger incentives for non-ETS firms, is therefore possible that policy aimed at non-ETS firms has been more efficient in providing abatement than the EU ETS. If a 'tight' cap has been in place, the ETS has ensured that abatement has occurred at a lower MAC, but does not speak to the effectiveness of the ETS in delivering abatement.

8 Conclusion

This dissertation has aimed to investigate the causal impact of EU ETS regulation on Norwegian manufacturing, mining and quarrying industries. Economic theory on environmental taxation and cap-and-trade predicts that imposing a cost of pollution leads to abatement, and that the price of allowances is determined by the marginal abatement cost of participating firms. I analyse the impact relative to the counterfactual scenario of non-ETS firms, implicitly assuming that regulation such at the CO_2 tax would have applied to ETS firms as well in the absence of emissions trading. Using Norwegian microdata I find a negative impact of around 20–40 % in phase II, similar to the findings of Klemetsen, Rosendahl, and Jakobsen (2020) and Petrick and Wagner (2014). I find a positive impact in phase III of around 30–60 %. While the findings rest heavily on strict assumptions, they are consistent in sign and relatively consistent in magnitude across model specifications.

Unfortunately, these results cannot be said to hold for firms off the common support. This lack of generalisability is a key caveat of the methodologies applied. It follows from the nature of the data and the regulation, making it difficult to find comparable ETS and non-ETS firms within the same industry. It is a point worth making that this conclusion holds for many of the microdata based EU ETS papers. While different samples may see different industries represented on the common support, this is typically not specified in the papers. There may be no apparent reason why the industries on the common support differ fundamentally for industries off the common support, we also cannot argue that random selection causes these problems. This means that although it cannot be ruled out that the estimates for τ_{ATT} found in this dissertation – and other EU ETS papers – hold for industries beyond those used to estimate the models, the results cannot be generalised. This calls for further research applying a wider range of methodologies.

The consistently negative and significant estimates of the impact of phase II are explained, in part, by high EUA prices at the beginning of the phase and harvesting of 'low-hanging fruit', where the marginal cost of abatement was low. In phase III, I find consistently positive and significant estimates of the impact. I interpret the lack of abatement found as a result of either a non-binding cap on emissions in phase III, as suggested by the large excess allowance supply, or of diverse marginal abatement costs between countries and firms. The negative estimates likely rest on my counterfactual assumption – that the climate change policies aimed at non-ETS firms may have been more efficient at providing abatement within Norwegian borders.

This paper did not set out to provide evidence of the effectiveness of the EU ETS outside of Norway, nor has it done so. However, to the extent that Norwegian firms have contributed to EU wide abatement in the ETS sector, this has likely been through driving up the price of EUAs, triggering abatement in other countries. For some, this may be taken as evidence for the ETS succeeding in minimising the cost of the abatement necessary to avoid catastrophic climate change. For others, it may suggest that the oil-rich country Norway is leaving the job of cleaning up to the more developed countries. The debate continues and more evidence is needed to shed light on the processes through which the EU ETS impacts mitigation of climate change.

Bibliography

Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey Wooldridge. 2017. *When Should You Adjust Standard Errors for Clustering?* Working Paper, Working Paper Series 24003. National Bureau of Economic Research. doi:10.3386/w24003.

Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin. 1996. "Identification of Causal Effects Using Instrumental Variables." *Journal of the American Statistical Association* 91 (434): 444–455. doi:10.1080/01621459.1996.10476902.

Angrist, Joshua D., and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion.* Princeton/Oxford: Princeton University Press. ISBN: 978-0-691-12035-5.

Angrist, Joshua, and Jinyong Hahn. 2004. "When to Control for Covariates? Panel Asymptotics for Estimates of Treatment Effects." *The Review of Economics and Statistics* 86 (1): 58–72. doi:10.1162/003465304323023679.

Arellano, Manuel, and Stephen Bond. 1991. "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *The Review of Economic Studies* 58 (2): 277–297. doi:10.2307/2297968.

Baumol, William J., and Wallace E. Oates. 1988. *The Theory of Environmental Policy.* 2nd ed. Cambridge University Press. doi:10.1017/CBO9781139173513.

Bel, Germà, and Stephan Joseph. 2015. "Emission Abatement: Untangling the Impacts of the EU ETS and the Economic Crisis." *Energy Economics* 49:531–539. doi:10.1016/j.eneco.2015.03.014.

Brink, Corjan, Herman R. J. Vollebergh, and Edwin van der Werf. 2016. "Carbon Pricing in the EU: Evaluation of Different EU ETS Reform Options." *Energy Policy* 97:603–617. ISSN: 0301-4215. doi:10.1016/j.enpol.2016.07.023.

Commission, European. 1992. "Proposal for a Council Directive Introducing a Tax on Carbon Dioxide Emissions and Energy. COM (92) 226 Final, 30 June 1992." Accessed August 14, 2020. http://aei.pitt.edu/4830/.

Commission, European. 2020. "EU Emissions Trading System (EU ETS)." Accessed July 30. https://ec.europa.eu/clima/policies/ets_en.

Ellerman, A. Denny, and Barbara K. Buchner. 2008. "Over-Allocation or Abatement? A Preliminary Analysis of the EU ETS Based on the 2005–06 Emissions Data." *Environmental and Resource Economics* 41 (2): 267–287. doi:10.1007/s10640-008-9191-2.

Ellerman, A. Denny, Claudio Marcantonini, and Aleksandar Zaklan. 2016. "The European Union Emissions Trading System: Ten Years and Counting." *Review of Environmental Economics and Policy* 10 (1): 89–107. doi:10.1093/reep/rev014.

Forbes, Kevin F., and Ernest M. Zampelli. 2019. "Wind Energy, the Price of Carbon Allowances, and CO2 Emissions: Evidence from Ireland." *Energy Policy* 133:110871. doi:10.1016/j.enpol.2019.07.007.

Goulder, L. H., and I. W. H. Parry. 2008. "Instrument Choice in Environmental Policy." *Review of Environmental Economics and Policy* 2 (2): 152–174. doi:10.1093/reep/ren005.

Hintermann, Beat, Sonja Peterson, and Wilfried 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 (1): 108–128. doi:10.1093/reep/rev015.

Holzer, Harry J., Richard N. Block, Marcus Cheatham, and Jack H. Knott. 1993. "Are Training Subsidies for Firms Effective? The Michigan Experience." *Industrial and Labor Relations Review* 46 (4): 625–636. doi:10.2307/2524308.

Klemetsen, Marit, Knut Einar Rosendahl, and Anja Lund Jakobsen. 2020. "The Impacts of the EU ETS on Norwegian Plants' Environmental and Economic Performance." *Climate Change Economics (CCE)* 11 (01): 1–32. Accessed June 5, 2020. https://ideas.repec.org/a/wsi/ccexxx/v11y2020i01ns2010007820500062.html.

Löschel, Andreas, Benjamin Johannes Lutz, and Shunsuke Managi. 2019. "The Impacts of the EU ETS on Efficiency and Economic Performance – An Empirical Analyses for German Manufacturing Firms." *Resource and Energy Economics,* Recent Advances in the Economic Analysis of Energy Demand - Insights for Industries and Households, 56:71–95. doi:10.1016/j.reseneeco.2018.03.001.

Martin, Ralf, Mirabelle Muûls, Laure B. de Preux, and Ulrich J. Wagner. 2014. "On the Empirical Content of Carbon Leakage Criteria in the EU Emissions Trading Scheme." *Ecological Economics* 105:78–88. doi:10.1016/j.ecolecon.2014.05.010.

Martin, Ralf, Mirabelle Muûls, and Ulrich J. Wagner. 2016. "The Impact of the European Union Emissions Trading Scheme on Regulated Firms: What Is the Evidence after Ten Years?" *Review of Environmental Economics and Policy* 10 (1): 129–148. doi:10.1093/reep/rev016.

NEA. 2010. *Klimakur 2020. Tiltak og virkemidler for å nå norske klimamål mot 2020 - Miljødirektoratet* TA 2590/2010. Climate and Pollution Agency (Klima- og forurensningsdirektoratet). Accessed September 13, 2020. https://www.miljodirektoratet.no/publikasjoner/publikasjoner-fra-klif/2010/mars/klimakur-2020.--tiltak-og-virkemidler-for-a-na-norske-klimamal-mot-2020/.

———. 2020. *Klimakur 2030: Tiltak og virkemidler mot 2030* M-1625. Norwegian Environmental Agency (Miljødirektoratet). Accessed August 26, 2020. https://www.miljodirektoratet.no/globalassets/publikasjoner/m1625/m1625.pdf.

Petrick, Sebastian, and Ulrich J. Wagner. 2014. *The Impact of Carbon Trading on Industry: Evidence from German Manufacturing Firms.* SSRN Scholarly Paper ID 2389800. Rochester, NY: Social Science Research Network. Accessed March 8, 2020. https://papers.ssrn.com/abstract=2389800.

Pigou, Arthur Cecil. 1924. *The Economics of Welfare.* 2nd ed. London: Macmillan.

Rosendahl, Knut Einar, and Halvor Briseid Storrøsten. 2011. *Output-Based Allocation and Investment in Clean Technologies.* Working Paper 644. Discussion Papers. Accessed September 22, 2020. https://www.econstor.eu/handle/10419/192626.

Teixidó, Jordi, Stefano F. Verde, and Francesco Nicolli. 2019. "The Impact of the EU Emissions Trading System on Low-Carbon Technological Change: The Empirical Evidence." *Ecological Economics* 164. doi:10.1016/j.ecolecon.2019.06.002.

Union, European. 2015. *EU ETS Handbook.* Accessed July 1, 2020. https://ec.europa.eu/clima/sites/clima/files/docs/ets_handbook_en.pdf.

Vona, Francesco. 2019. "Job Losses and Political Acceptability of Climate Policies: Why the 'Job-Killing' Argument Is so Persistent and How to Overturn It." *Climate Policy* 19 (4): 524–532. doi:10.1080/14693062.2018.1532871.

Wagner, Ulrich J., Dana Kassem, Andreas Gerster, Jurate Jaraite, Marit E. Klemetsen, Marita Laukkanen, Ralf Martin, et al. 2020. *Carbon Footprints of European Manufacturing Jobs: Stylized Facts and Implications for Climate Policy.* SSRN Scholarly Paper ID 3577550. Rochester, NY: Social Science Research Network. doi:10.2139/ssrn.3577550.

Wagner, Ulrich J., Mirabelle Muûls, Ralf Martin, and Jonathan Colmer. 2014. "The Causal Effects of the European Union Emissions Trading Scheme: Evidence from French Manufacturing Plants [Working Paper]." Accessed September 19, 2020. http://conference.iza.org/conference_files/EnvEmpl2014/martin_r7617.pdf.

Weitzman, Martin L. 1974. "Prices vs. Quantities." *The Review of Economic Studies* 41 (4): 477–491. doi:10.2307/2296698.

Wettestad, Jørgen, and Torbjørg Jevnaker. 2016. *Rescuing EU Emissions Trading: The Climate Policy Flagship.* London: Palgrave Macmillan UK. doi:10.1057/978-1-137-56674-4_6.

———. 2018. "EU Emissions Trading: Frontrunner - and 'Black Sheep'?" In *The Evolution of Carbon Markets: Design and Diffusion,* edited by Jørgen Wettestad and Lars H. Gulbrandsen. London: Routledge.

Wooldridge, Jeffrey M. 2010. *Econometric Analysis of Cross Section and Panel Data.* The MIT Press. doi:10.2307/j.ctt5hhcfr.

Theory

Externalities and pricing emissions

The idea of externalities and pricing them in order to internalise the social cost of production dates back to Pigou (1924). In general, externalities occur when a cost (or benefit) is imposed upon someone in a way that is not taken into account by the agent who generates the externality. Baumol and Oates (1988) establish two conditions:

a) The utility or production relationship of some individual A include non-monetary variables chosen by someone else who is not considering the effect on A's welfare

b) The decision maker does not pay a price (or receive compensation) equivalent to the effect their activity has on the welfare of A

These conditions are intuitive: If an agent's activity has no impact on others, it should be no concern of theirs. If the impact they have on others is captured through pricing in the market, marginal cost and marginal willingness to pay should equalise and generate the socially optimal production level. Conditions a) and b) are therefore key to solving externalities. This limits my discussion to market failures, and does, for instance, not treat properties of (perfect) competition as externalities. The conditions also exclude other market failures, such as increasing returns to scale (typically warranting a "natural monopoly") (Baumol and Oates 1988). The externalities here are therefore the ones that necessarily prevent the economy from reaching a Pareto efficient allocation.

Greenhouse gas emissions, which are produced by private firms and consumers, increase global temperatures and reduce the utility achieved by all from the public good of a stable climate. The emissions do not incur any cost to the emitter without policy intervention. It is therefore a clear example of a negative externality.

The idea of setting a "price on carbon" arises from condition b). While it may be impossible (or at least, very inefficient) to change the inherit relationship of an agent imposing a cost (or benefit) on others, the externality problem can be resolved by "internalising" this cost (or benefit). A "Pigouvian tax" sets the tax equal to the marginal cost of abatement, such that the socially optimal level of pollution is reached. Quantity controls, particularly cap-and-trade systems, have gained increasing popularity among politicians and economists alike for their argued efficiency and flexibility. In an economy of perfect information, the concept of pricing emissions is illustrated in figure 1.

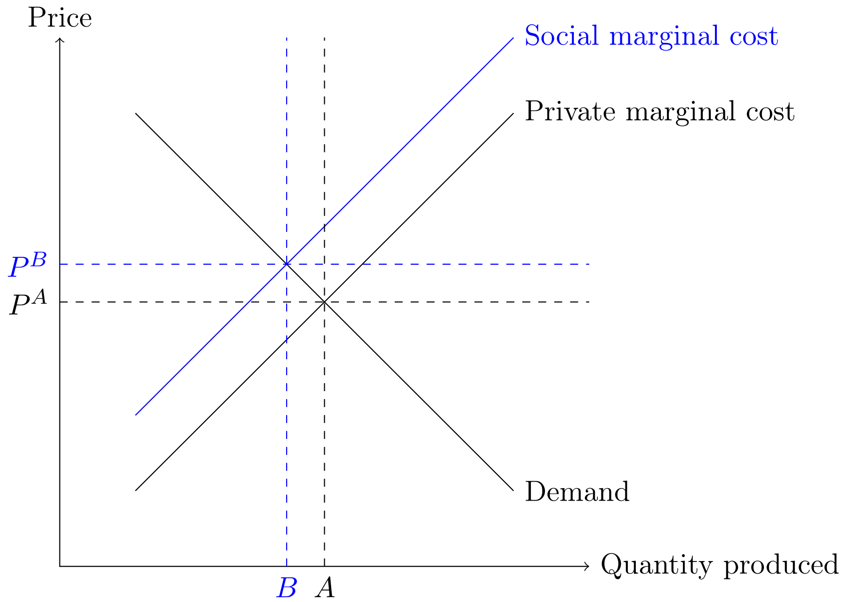

Figure 1: Externality

Figure 1 is a standard market equilibrium of car driving. The line "Private marginal cost" is the private marginal cost of driving, e.g. fuel costs, car rental, etc. The blue line illustrates the social marginal cost of carbon, which equals private marginal cost plus the cost of the externality caused by greenhouse gas (GHG) emissions. In absence of policy, the equilibrium quantity is given by A at market price P^A. The optimal quantity would be B at optimal price P^B. This diagram illustrates how externalities cause market prices to lie below the optimal price and cause overconsumption. A Pigouvian tax is given by $\tau = P^B - P^A$ and thereby shifts the market quantity from A to the optimal point B.

Emissions trading

Emissions trading systems, or cap-and-trade systems, issue permits that cap the amount of emissions within a given period. These permits can then be traded, ensuring that abatement occurs where the marginal cost is lowest, thereby achieving the lowest possible marginal cost of abatement. The instrument is less suited if pollution is local, but for global externalities such as greenhouse gases, it is irrelevant where pollution occurs or is abated. The optimal cap in the cap-and-trade system yields the market optimum in figure 1 by setting the cap equal to B. The new market equilibrium is then achieved at price P^B.

At first glance, Pigouvian taxes and cap-and-trade seem very similar systems. However, economies are rarely as simple as in figure 1. Pigouvian taxes and cap-and-trade are equivalent only under complete information and in the absence of any other market failures. If there is uncertainty about abatement costs quantitative controls ensure that the targeted emissions level is reached, while price controls ensure the targeted marginal abatement cost is reached. If the marginal damage function is inelastic, quantitative controls are preferred (Weitzman 1974).

While taxes are vulnerable to the level of the tax set, quantitative controls are equally vulnerable to the quantity of permits set. The key criticism against the EU ETS has been its very high cap on emissions, leading to a large surplus in the market for allowances. This is likely to have severely limited the impact of the ETS on abatement.

Further, cap-and-trade systems have gained popularity due to being perceived as more politically feasible. The EU, for instance, does not have a common fiscal policy and the Commission was therefore unsuccessful in proposing an EU-wide carbon tax (see Commission 1992). The importance of feasibility in implementation of policy instruments should not be ignored (Vona 2019).

Auctioning or grandfathering

Emissions trading also raises the question of how permits should be issued. In the EU ETS, permits are issued for free either on the basis of historical emissions, so-called grandfathering, or auctioned. Regardless of the allocation method the market price of allowances should equal the marginal cost of the last unit abated – as long as the cap is tight (i.e. does not exceed total emissions) and the market is competitive. Whether allowances are allocated for free or auctioned off, firms will only be willing to buy allowances if the price of doing so is lower than the marginal cost of reducing emissions by one tonne. The primary theoretical consideration in the choice of allocation method is therefore the trade-off between raising revenue and imposing a cost on firms.

Auctioning of permits increases government revenue and induces a revenue recycling effect (Goulder and Parry 2008). This could allow governments to reduce other distortionary taxes, e.g. in the labour market, adding an additional efficiency gain to that of reducing the externality. By contributing to governmental income auctioning revenue can therefore improve the overall efficiency of tax systems[1]. In general, theoretical models of emission trading assume that the same allocation method is chosen for industries. As described in section 2.3, allocation has not been homogenous across firms or industries.

1. This paper cannot capture such general equilibrium effects, but the reader should be aware that the introduction of more auctioning may have had additional efficiency gains. However, auctioning imposes a cost on firms and may increase the risk of "carbon leakage", i.e. production simply being moved outside the EU, indicating other potential negative general equilibrium effects. Computable general equilibrium (CGE) analyses of the EU have been conducted (see e.g. Brink, Vollebergh, and van der Werf 2016), but are outside the scope of this paper.

Phases and development of the EU ETS

This appendix sets out the structural design of the EU ETS and its development.

Early days: Phase I

The EU ETS was first adopted as EU law in 2003, following the targets set for the EU in the 1997 Kyoto Protocol (Wettestad and Jevnaker 2016). Phase I, the "pilot phase", span from 2005 through 2007 and was characterised by decentralisation. The main target industries were power generation and energy-intensive industries. Capacity thresholds ensured that only those plants exceeding a certain thermal input or production capacity were regulated. Research and development activities were excluded. In particular, this means that energy generation, oil refineries, production and processing of ferrous metals, mineral industry (e.g. cement, glass, ceramics) and pulp and paper were regulated (Directive 2003/87/EC).

There was no central cap and allowances were generally distributed for free as determined by member states. At least 90 % of allowances had to be grandfathered due to concerns about imposing competitive disadvantages on European firms in the global market (Wettestad and Jevnaker 2018). Phase I saw different strategies of implementation in different states and volatile allowance prices (Wettestad and Jevnaker 2016). Firms were allowed to bank (save) allowances for future years, but only within the first phase. Flexibility mechanisms from the Kyoto Protocol were put in place, allowing firms to buy Clean Development (CDM) allowances instead of EUAs. This essentially allowed EU firms to pay for abatement by firms outside the EU in addition to the EU ETS cap. Phase I saw initially high allowance prices, overallocation of allowances and eventually a near-zero carbon price (Wettestad and Jevnaker 2016).

The EU ETS put to the test: Phase II

Phase II followed the Kyoto commitment period of 2008–2012. Coinciding with the 2008/2009 financial crisis and following recession, the new institution was put to the test. Free allowances were still the main rule, with 10 % being auctioned, but fewer allowances were introduced in the market. External credits (CDM and Joint Implementation, JI) could still be used. Allowances could now be banked for future phases in an aim to stabilise prices over time. The regulated firms were mainly the same as in phase I, but aviation was included from 2012.

Allowance prices fell dramatically as a large surplus emerged throughout the phase. By the end of 2012 the allowance price had fallen below 10 euros. At this point, criticism of the EU ETS had reached high levels. Climate activists considered the low allowance prices and questioned its ability to deliver abatement. Others were sceptical of imposing additional costs on firms already struggling from the aftermath of the recession. As a consequence, controversy was prominent over reforming the ETS and the situation seemed gridlocked (Wettestad and Jevnaker 2016).

The Commission proposed a new reform in 2012, including measures to tighten the cap, reduce the surplus and manage the allowance price.

This controversy is important in understanding the adjustment of firms to the EU ETS. With low carbon prices and a gridlocked EU, it was natural to question the extent to which the ETS would interfere with the "business as usual" scenario. These expectations may have influenced the willingness of firms to introduce costly abatement measures or invest in "green" innovation.

Refinement and renewed hope: Phase III

In 2008, the second EU ETS directive (Directive 2008/101/EC) was adopted, introducing the first major reform of the EU ETS (Wettestad and Jevnaker 2018). Phase III (2013–2020) introduced a paradigm shift for the EU ETS, introducing far more centralisation. A single, EU-wide cap was introduced and allocation of permits became the rule rather than the exception. Emissions are to be reduced by 21 % from 2005 levels by 2020, meaning a linear reduction of 1.74 % per year.

Allowances were to be allocated according to a fully harmonised system, not national rules and no new external credits could enter the system (Wettestad and Jevnaker 2016). While only CO_2 emissions were regulated in phases I and II, N_2O and PFCs were included from 2013. The regulated industries remained largely the same, but was slightly increased (e.g. by including aluminum production). Electricity-generating industry now faced 100 % auctioning, while 20 % (increasing to 70 % by 2020) were auctioned to other industries. Some particularly exposed industries were allocated free allowances for the duration of the phase, with technology-based benchmarks replacing historical emissions as allocation criteria (Wettestad and Jevnaker 2016).

A surplus of allowances banked in phase II remained as phase III began. In response to the decline in EUA prices in phase II, a backloading mechanism was introduced. This temporary mechanism ensured that releasing new allowances to the market would be postponed if the surplus was too large. In 2015 the EU agreed upon a market stability reserve (MRS) and a "complete ETS overhaul", a far more complex and comprehensive reform (Wettestad and Jevnaker 2016). The MRS releases new allowances to the market if the supply is very low, and withdraws some if the surplus is very high. As such, the MRS aims to reduce fluctuations in the market. Further, its introduction formalises a new dimension of the EU ETS – if abatement exceeds the target, it does not necessarily lead to increased emissions elsewhere in the EU or in the future. While the MRS has only been in place for a little more than a year, its announcement in 2015 may have lead to anticipation effects among forward-looking firms.

Long-term outlooks: Phase IV

In phase IV (2021-2030), the centralisation and tightening of the cap is expected to continue. Phase IV outlooks are relevant to this paper to the extent that anticipation effects can be expected to influence firm behaviour. The MRS will continue to play an important role, and allowances will be reduced by 2.2 % per year rather than 1.74 %. Free allocation will be limited to the sectors deemed at the highest risk of relocation and new or growing installations (Commission 2020). New funds will be put in place to support innovation in general and lower-income member states in particular. In sum, this points towards a more comprehensive and "tight" market, possibly causing firms to expect allowance prices to continue to rise.

Norway and emissions trading

Norway only joined the EU ETS through the EEA in 2008, when phase II was introduced. Norway instead prepared for EU ETS inclusion by introducing its own "phase I" emissions trading system from 2005. While not formally linked to the EU ETS, EU ETS allowances (EUAs) could, among others, be submitted if firms exceeded their initial emissions allowances (Klemetsen, Rosendahl, and Jakobsen 2020). The regulation was designed to ensure that the participation criteria were the same as in the EU ETS. Around 10 % of emissions were regulated[2].

The regulated status of firms in the first and second phases were determined by the base period 1998–2001 (with some adjustments made for firms that started or changed their production after 2001). When Norway introduced its national emissions trading system (ETS) in 2005 its intention was to lay the groundwork for future close links with the EU ETS.

At the end of phase III, the majority of emissions from Norwegian manufacturing firms are regulated by the EU ETS (NEA 2020). The non-regulated firms are generally smaller with far lower emissions than the regulated firms.

2. Allowances were allocated based on plants' emissions in 1998–2001.

The CO_2 tax: Email exchange with the Ministry of Finance

Available upon request.