

# 1925

Discussion  
Papers

## Coase and Cap-and-Trade: Evidence on the Independence Property from the European Carbon Market

Opinions expressed in this paper are those of the author(s) and do not necessarily reflect views of the institute.

#### IMPRESSUM

© DIW Berlin, 2021

DIW Berlin  
German Institute for Economic Research  
Mohrenstr. 58  
10117 Berlin

Tel. +49 (30) 897 89-0  
Fax +49 (30) 897 89-200  
<http://www.diw.de>

ISSN electronic edition 1619-4535

Papers can be downloaded free of charge from the DIW Berlin website:  
<http://www.diw.de/discussionpapers>

Discussion Papers of DIW Berlin are indexed in RePEc and SSRN:  
<http://ideas.repec.org/s/diw/diwwpp.html>  
<http://www.ssrn.com/link/DIW-Berlin-German-Inst-Econ-Res.html>

# Coase and Cap-and-Trade: Evidence on the Independence Property from the European Carbon Market

Aleksandar Zaklan\*

January 2021

Updated version of DP 1850. Please cite this updated version.

This paper tests the independence property under the Coase Theorem in a large multinational cap-and-trade scheme for greenhouse gas emissions, the EU Emissions Trading System (EU ETS). I analyze whether emissions of power producers regulated under the EU ETS are independent from allowance allocations, leveraging a change in allocation policy for a difference-in-differences strategy. The evidence suggests that the independence property holds overall and for larger emitters. It fails for small emitters, indicating that transaction costs distort their emission decisions. However, due to their small share of aggregate emissions the independence property remains intact at the sector level.

**JEL codes:** Q54, Q58, Q52, L94

**Keywords:** Coase theorem, independence property, cap-and-trade, EU ETS, greenhouse gas emissions.

---

\*DIW Berlin. Mohrenstr. 58, 10117 Berlin, Germany. E-Mail: azaklan@diw.de. Ph. +49-30-89789-515. I especially thank my Jean Monnet advisor Denny Ellerman, as well as Jan Abrell, Jochen Diekmann, Mirjam Kosch, Jan Marcus, Sebastian Pieper, Brendan Restrepo, Luis Sarmiento, Sebastian Schwenen and Ulrich Wagner for helpful discussions. I also thank participants at the Verein für Socialpolitik Annual Conference, EAERE Annual Conference, Mannheim Energy Conference, EUI Climate Conference, Atlantic Workshop on Energy and Environmental Economics, and CERE Research Seminar at Umea University for useful comments. All remaining errors are my own.

# 1. Introduction

Cap-and-trade has become an increasingly popular approach to mitigating global and local environmental externalities. From its origins in the U.S. it expanded to a large number of countries and is seen as an option for global greenhouse gas control under the Paris Agreement (UNFCCC, 2015).<sup>1</sup> Its attractiveness rests on its promise of cost-effectiveness: By combining a cap with the trading mechanism abatement takes place where it is cheapest, equalizing marginal abatement costs across regulated units in equilibrium and minimizing aggregate costs for a given amount of abatement. Based on the seminal insight by Coase (1960) and extended to cap-and-trade by Montgomery (1972), under certain conditions agents' equilibrium solutions for the amount of the externality produced are independent from the allocation of allowances. Changes in allocations are limited to having distributional effects. Dependence between allocations and emissions would be a symptom of underlying friction in the market. In the classic Coasean case it indicates the presence of transaction costs (Coase, 1960; Stavins, 1995).<sup>2</sup>

The independence property is important in real-world contexts, as policy makers frequently use free allocation in the political bargaining process to establish and/or maintain cap-and-trade schemes. For example, in the EU Emissions Trading System (EU ETS) most of the cap was freely allocated during its early years, and close to half of the cap is still distributed for free today to mitigate carbon leakage risk, despite concerns that emitters may be overcompensated (Hepburn et al., 2013; Martin et al., 2014). A cap-and-trade scheme in which emissions are independent from allocations is attractive for policy makers, as they can focus on whether free allocation is desirable from an equity perspective (e.g. Ambec and Ehlers, 2016; Fleurbaey and Maniquet, 2018). Empirical evidence on whether the independence property holds in operational cap-and-trade schemes is therefore highly relevant to shed light on whether a liberal use of free allocation should be avoided. Yet despite its importance the existing empirical evidence is extremely scarce. To my knowledge, the existing peer-reviewed literature is limited to a contribution by Fowlie and Perloff (2013) on a scheme for air pollution in California; the only existing

---

<sup>1</sup>See, e.g., <https://icapcarbonaction.com/en/ets-map> for an overview of current cap-and-trade programs.

<sup>2</sup>Independence may also be affected by imperfect competition in the allowance market (Hahn, 1984; Liski and Montero, 2010), behavioral disparities between agents' willingness to pay and the willingness to accept (Kahneman et al., 1990; Hanemann, 1991). A further reason for a dependence between allocation and emissions can be capacity-based subsidies to mitigate the risk of carbon leakage (Löfgren et al., 2018) or incentives for firms to distort the allowance market due to the updating of free allocation (Harstad and Eskeland, 2010).

contribution on climate policy is a working paper by Reguant and Ellerman (2008). This paper therefore contributes to filling an important gap in the literature.

This analysis tests whether emissions are independent from allowance allocations for firms regulated under the world’s largest multilateral cap-and-trade scheme, the EU ETS. The EU ETS became operational in 2005 and covers most of the emissions from the European electricity supply and a number of manufacturing sectors. The analysis focuses on power producers, responsible for more than half of the EU ETS’s greenhouse gas emissions and arguably the most sophisticated sector under the EU ETS with respect to the allowance trade.<sup>3</sup> However, not all producers may be able to participate in the allowance trade equally. Smaller emitters face relatively higher transaction costs due to resource constraints. E.g., as they often do not have the internal capacity to directly participate in allowance auctions, they need to use intermediaries and pay additional per-unit trading fees (Eden et al., 2019), while larger emitters can more easily reap economies of scale (Jaraitė et al., 2010; Heindl, 2017). These constraints introduce a significant wedge in average transaction costs between small and large emitters. Using data for 2005-2007 Jaraitė et al. (2010) estimate average transaction costs of about 2 euro per ton of CO<sub>2</sub> for small emitters and 0.05 euro per ton for large emitters. Based on survey data from German companies Heindl (2012) estimates that average transaction costs were up to 1 euro per ton of CO<sub>2</sub> for small emitters and quickly declined for larger emitters. Transaction costs of this magnitude would represent a significant share of the allowance price, which was in the range of 5-8 euro per ton during the majority of the sample period in this paper (Löfgren et al., 2018). The problem of transaction costs for small emitters has also been recognized by policy makers. Current EU ETS rules include a (little-used) opt-out provision for small emitters in exchange for national-level ”equivalent measures” (Article 27, EU, 2009, 2018). It is therefore plausible that at least some producers exhibit dependence between allocations and emissions, which may affect independence at the sector level. As the EU ETS has become a benchmark for other cap-and-trade initiatives around the world targeting greenhouse gas emissions, such an analysis is also of interest to other cap-and-trade programs looking at the EU ETS as a blueprint.

The regression analysis pursues a difference-in-differences (DID) research design using

---

<sup>3</sup>Fabra and Reguant (2014) find that electricity producers fully internalize the opportunity cost of emissions.

a panel of power producers from coal and gas during the period 2009 to 2017.<sup>4</sup> The causal effect of allowance allocations on installations' emissions is identified by leveraging a change in allocation policy: In 2013, in the majority of EU member states power producers lost a large share of their freely allocated allowances. Since then they must purchase them through auctions, either directly or through intermediaries, or on the secondary market. However, this decrease in allocations did not apply to all electricity generators in the same manner. While power generators generally faced full auctioning starting in 2013, a special provision – the so-called 10c rule – preserved a significant amount of free allocation for producers in eight EU member states. The analysis thus compares emissions decisions of producers who lost most of their free allocation (treatment group) with those who continue receiving substantial free allocation under the 10c rule (control group).<sup>5</sup>

The divergence in allocation levels and thus in the need to purchase additional allowances provides substantial variation in total compliance costs across treatment and control groups. I hypothesize that if transaction costs depend on the amount of allowances purchased, the divergence in allocation levels will expose treated producers to a differential in transaction costs. If treated producers experience greater transaction costs, they will purchase fewer allowances than socially optimal and instead abate more than socially optimal. In this case, their emissions will decrease compared to those of the control group. This problem is likely to be more pronounced for small emitters.

In addition to investigating the independence property using the pooled firm-level sample, I also distinguish between small and larger emitters, to test whether a possible dependence between allocations and emissions may be masked in the full sample. Moreover, I use disaggregated installation-level data to further probe the analysis at the firm level. Installations are further subdivided by fuel type and owner type.

It is important to control for sample selection, which occurs at the country level, as this is where the 10c rule was negotiated. I argue that political preferences for continued free allocation are unobservable but fixed, so that it is important to exploit the panel nature of the data and control for producer fixed effects. Furthermore, I include the main country-level electricity market outcomes relevant to the production (and thus emission) decision of a fossil fuel plant. The analysis also controls for cross-country drivers of emissions,

---

<sup>4</sup>Oil-fired facilities are excluded from the analysis as oil, first, does not play a large role in European power generation and, second, oil-fired generation hardly plays a role in the control group.

<sup>5</sup>UK producers are excluded from the analysis, as the UK Carbon Price Support, a carbon tax on top of the allowance price for electricity generators, was also introduced in 2013 (cf. e.g. Abrell et al., 2019).

either implicitly through year fixed effects, or explicitly, by including allowance and fuel prices. Estimations are run on full and matched samples.

The regression analysis fails to reject the independence property for the power sector as a whole and for larger emitters representing the lion's share of power sector emissions. For these firms a change in the level of free allocation does not cause a change in emissions. However, the independence property fails for small emitters, for whom I estimate a significant negative relationship between allowance allocations and emissions. The change in allocation levels is estimated to lead to a decrease in the emissions of this group of firms by about 20%. This negative relationship is in line with the prediction from theory that in the presence of transaction costs producers not receiving enough free allocation to cover all their emissions abate more than would be optimal in a transaction-cost-free setting. I thus conclude that there is evidence of transaction costs affecting emission decisions by small emitters. However, as they represent a small share of the power sector's total emissions, distorted emission decisions by this group of firms do not significantly affect the rest of the power sector.

The findings in this paper are policy relevant, as a further expansion of cap-and-trade is considered to be one option for international greenhouse gas control. My analysis suggests that policy makers may use free allocation as a tool in the political bargaining process without significantly distorting the cap-and-trade program, as long as the setting is comparable to the one in this paper. However, my result for small emitters suggests that further economizing on the cost of market participation may increase market participation and political acceptance of emissions trading. Two caveats are in order with respect to this policy conclusion: First, even though allocations and emissions are found to be independent *ex post* in this study, when deciding to liberally use free allocation policy makers should focus on minimizing distortions to the allowance market due to incentive effects entering firms' decisions through the dynamics of free allocation. As long as current decisions by firms are tied to future allocation levels, firms have an incentive to distort the allowance market to secure higher allocations in the future (e.g. Boehringer and Lange, 2005; Rosendahl, 2008; Harstad and Eskeland, 2010; Meunier et al., 2018). Second, even if the free allocation scheme is designed to minimize distortions, how the cap is distributed has important implications for environmental justice.

To the best of my knowledge, this paper is the first to comprehensively analyze the independence property in a large-scale multinational cap-and-trade program for greenhouse

gas emissions using a quasi-experimental approach. The focus on emitter size is also novel. In terms of empirical strategy, this paper is most similar to Fowlie and Perloff (2013), who consider the independence property for nitrogen oxide ( $\text{NO}_x$ ) emissions under the Regional Clean Air Incentives Market (RECLAIM), a cap-and-trade scheme in Southern California covering emissions of local air pollutants  $\text{NO}_x$  and sulfur oxides ( $\text{SO}_x$ ). In contrast, this paper focuses on a large-scale program targeting greenhouse gas emissions, a pollutant without a commercially viable end-of-pipe abatement option.<sup>6</sup> Reguant and Ellerman (2008) address the independence property for Spanish power plants during the pilot phase of the EU ETS using a short-run power production model. Compared to Reguant and Ellerman (2008) this paper mainly contributes by using quasi-experimental variation to identify the effect of allocation on emissions, by expanding the analysis beyond Spanish coal plants, beyond the EU ETS pilot phase dominated by free allocation and an inability to bank allowances, and by paying attention to the role of small emitters.<sup>7</sup> My finding that the independence property fails for small emitters adds a new dimension to the existing literature on the independence property in cap-and-trade schemes. At the same time, the conclusions from this paper are in line with the main results from the existing literature that finds support for the independence property in the overall allowance market.

In another related contribution, Montero et al. (2002) show that transaction costs can be high if an allowance market does not have a strong institutional foundation. The analysis in this paper suggests that the design of the EU ETS led to relatively low transaction costs for the majority of producers. A further related but distinct contribution is Carlson et al. (2000), who estimate abatement cost curves for the U.S.  $\text{SO}_2$  trading program and evaluate the extent to which cost economies were reaped during its early years. The analysis in Carlson et al. (2000) is complementary to this paper, as it directly focuses on the ability of a cap-and-trade scheme to deliver on its promise of exploiting potential gains in trade compared to less flexible schemes.

This paper is also related to the literature analyzing the effectiveness of cap-and-trade programs, and especially the effectiveness of the EU ETS. Bayer and Aklin (2020) use sector-level data to estimate the effectiveness of the EU ETS with respect to emission abatement, while Colmer et al. (2020) and Petrick and Wagner (2014) estimate the effec-

---

<sup>6</sup>Carbon capture and storage (CCS) is one potential option. However, it is currently not commercially viable and may or may not become viable in the future.

<sup>7</sup>See, e.g., Ellerman et al. (2016) for a description of the progression of the EU ETS through its different trading periods including the main rule changes.



tiveness of the EU ETS with respect to abatement of French and German manufacturing firms, respectively. Dechezleprêtre et al. (2019) and Naegele and Zaklan (2019) are examples of contributions analyzing its effectiveness with regard to international emission abatement, i.e. whether abatement in Europe is negated by carbon leakage. Cael and Dechezlepretre (2016) and Cael (2020) analyze the EU ETS's impact on firm-level innovation activity, while Abrell et al. (2019) evaluate the effectiveness of the UK Carbon Price Support, a carbon tax introduced for UK power producers on top of the EU ETS allowance price. The conclusions in this paper are relevant to the question whether the effectiveness of the EU ETS identified in this literature was achieved at greater than minimal aggregate cost due to transaction costs.

The remainder of the paper proceeds as follows: Section 2 provides background on free allocation for electricity generators in the EU ETS. Section 3 presents the data, including the matching and descriptive evidence on parallel trends, while Section 4 presents the research design. Sections 5 and 6 contain the empirical results, and Section 7 summarizes and concludes.

## **2. Background: Allocation of allowances to power producers**

The EU Directive 2009/29/EC (EU, 2009) provides the legal framework for the EU ETS during Phase III, covering the period 2013-2020. Article 10 determines that power generators – in general – shall not receive any free allocation starting in 2013.<sup>8</sup> Instead, they had to purchase allowances either at auctions or on the secondary market, increasing auctioning revenues for EU member states.<sup>9</sup> The change in allocation policy intended to eliminate windfall profits to power producers, who can fully pass through the emission costs to electricity prices (Fabra and Reguant, 2014). This policy change led to the majority of the European electricity sector transitioning rapidly from widespread free allocation to mostly full auctioning between 2012 and 2013 (see also Section 4.1). This group of producers constitutes the treatment group in the empirical analysis of this paper. Importantly, the loss in allocation did not differentiate by emitter size. Small producers were affected in the same way as large conglomerates with specialized trading departments.

---

<sup>8</sup>Cogeneration facilities simultaneously providing heating and cooling services could continue receiving free allocation for the emissions involved in generating heating or cooling.

<sup>9</sup>Based on available information auctioning revenues were mainly used to support renewable energies and/or energy efficiency measures. However, information on the use of auctioning revenues is incomplete. Whether the use of auctioning proceeds is additional to other country-level expenditures is also unclear (Löfgren et al., 2018).

In contrast, power generators from a select group of EU member states continued being eligible to receive free allocation after 2013 under Article 10c of Directive 2009/29/EC (EU, 2009).<sup>10</sup> They constitute the control group in the empirical analysis of this paper. Member states' governments could apply for an extension of free allocation if (a) their national electricity network was either not or only poorly connected to the interconnected European electrical grid or (b) if more than 30% of the country's electricity was generated from a single fossil fuel and GDP per capita was at most 50% of the EU level in 2006. Only power plants in existence or *physically initiated* by December 31, 2008 were eligible to receive allowances for free, to avoid new entries due to the allocation subsidy. Ten EU member states were eligible to provide free allocation to their power generators under the 10c rule, all of them countries who joined the EU after 2004. The countries eligible are Bulgaria, Cyprus, Czech Republic, Estonia, Hungary, Latvia, Lithuania, Malta, Poland and Romania. Out of the member states eligible for continued free allocation Latvia and Malta declined to use this option.

The stated goal of continued free allocation under Article 10c is to support producers in eligible member states with modernization measures, to avoid purely perpetuating windfall profits from free allocation. Allowances allocated for free under the 10c rule are subtracted from each member state's auctioning contingent, i.e. countries opting for free allocation under the 10c rule incurred losses in auctioning revenue. However, at the time when the 10c rule went into effect allowance prices were historically low, reducing the opportunity cost of continued transfers to domestic power producers through free allocation for the governments involved.

Each member state using the provision under the 10c rule had to submit a list of investment options for approval by the European Commission by September 30, 2011. The total value of investments had to be at least equal to the market value of freely allocated allowances at that time. Member states had to report to the Commission on the progress of investments each year. Stakeholders criticized the process under the 10c rule as lacking transparency. Evidence from the grey literature suggests that investment projects under the 10c rule mainly subsidized the continued production of fossil-fueled power by incumbent companies, and many proposed projects were considered doubtful with respect to emission reductions (CMW, 2016; Popa and Ceperaga, 2012).

As further elaborated on in Section 4.2, power generators from countries that pursued

---

<sup>10</sup>While Article 10c lays out the general rules for continued free allocation to power generators during Phase III, Communication 2011/C 99/03 (EU, 2011) fleshes out the implementation of Article 10c.

free allocation using the 10c rule constitute a valid control group. Possible deviations in emission patterns, e.g. due to free allocation being tied to the modernization measures, would bias the DiD analysis against finding a treatment effect, allowing for a lower-bound interpretation of the results.

## 3. Data

### 3.1. Data sources

Data on installations from the category 35 "Electricity, gas, steam and air conditioning supply" according to the two-digit NACE classification are obtained from the European Union Transaction Log (EUTL). The mapping of installations to sectors is provided by the European Commission and obtained through the Ownership Links and Enhanced EUTL Dataset Project (Jaraitė et al., 2016). I consider a balanced panel of installations reporting positive emissions for every year of the sample period and aggregate them to the firm level for the main regression analysis.<sup>11</sup> The sample period is 2009-2017.<sup>12</sup>

Note the UK Carbon Price Support (CPS), which introduced a minimum price for the emissions of electricity generators, was also instituted in 2013, the same year that the EU ETS introduced auctioning for electricity producers. To avoid attributing the significant abatement effects of the UK CPS (Abrell et al., 2019) to the loss in free allocation, I exclude installations located in the UK.

The main source of information on the main fuel type of installations is the Open Power System Data (OPSD) platform,<sup>13</sup>. Installations were matched by comparing installation names in the EUTL with names provided in the OPSD data. For some countries, coverage in the OPSD data was either unavailable or incomplete. In these cases, data from the Global Energy Observatory (GEO)<sup>14</sup> were used to determine the main fuel type, again by comparing EUTL installation names to facility names in the GEO data. For installations whose main fuel type could not be determined either using OPSD or GEO information a

---

<sup>11</sup>Focusing on the sample of "remainders" helps us avoid falsely associating the effects of plant shutdowns under the Large Combustion Plant Directive (LCPD) (EU, 2001) to the switch to auctioning in the EU ETS, as a number of coal-fired power plants ended operation due to the LCPD during our post-treatment period. Focusing on a balanced panel also provides a lower bound for the estimate of the treatment effect, if installations are closed due to the loss in free allocation.

<sup>12</sup>Observations from 2008, the year in which the financial crisis started, are excluded. I additionally exclude installations with extreme swings in emissions, where the minimum emissions are less than 5% of maximum emissions, e.g. due to maintenance shutdowns. Results are similar when other cut-off points are chosen.

<sup>13</sup><https://open-power-system-data.org/>

<sup>14</sup><http://globalenergyobservatory.org/>

I conducted a manual web search. Using this three-pronged approach I was able to assign the main fuel type to 88% of the installations identified as belonging to the power sector, representing about 97% of the emissions.

The country-level electricity market covariates electricity consumption, the amount of electricity produced from renewable energies and net electricity exports are from Eurostat. So are data on GDP. Data on allowance and fuel prices are from Bloomberg.

### 3.2. Descriptive overview

Table 1 provides a descriptive overview of the full estimation sample, by treatment and control groups. Note that despite the limited size of the control group the distributions of pre-treatment emissions among treated and control gas plants exhibit a good fit in terms of suitable matches (see Section 3.3) and parallel trends (see Section 3.4).

The top section of Table 1 presents the firm-level data, distinguishing the full sample of all firms, the small emitter subsample and the non-small emitter subsample. I follow the European Commission’s definition of a firm, denoted as ”account holders” in the EUTL data. Small emitters are defined as those below the 25th percentile of the distribution of firms’ 2009 emissions, while non-small emitters are the remaining firms, i.e. those at or above the 25th percentile of the distribution of firms’ 2009 emissions. Percentiles are calculated separately for the treatment and control groups, i.e. in the small emitter sample we compare the bottom 25% of emitters in the treatment group with the bottom 25% of emitters in the control group.<sup>15</sup>

Firms in the treatment group have fewer emissions, in terms of both the mean and the median, than those in the control group, reflecting the greater share of coal-fired power in the control group. The relative difference in emitter size is even more pronounced in the small firms subsample, while it is less stark in the non-small emitters subsample. In both groups and for most samples (with the exception of the small-emitter subsample), the distribution of emissions is skewed toward a small number of large emitters, as shown by the great difference between median and mean emissions. This is why emissions enter the regression analysis in logarithmic form.

The second block of Table 1 shows the installation-level data, first for all coal and gas

---

<sup>15</sup>Lowering the percentage cutoff further severely limits the sample size, especially in the control group.

Table 1: Descriptive overview

	Treatment group			Control group		
	Mean	Std. Dev.	Median	Mean	Std. Dev.	Median
<b>Firm-level data</b>						
<i>All firms</i>						
Emissions (tons of CO <sub>2</sub> )	1,742,58	6,848,563	191,182	2,890,725	7,183,806	682,723
No. firms	304			81		
No. observations	2,736			729		
<i>Small emitters</i>						
Emissions (tons of CO <sub>2</sub> )	13,795	11,453	10,293	105,150	67,940	91,114
No. firms	76			20		
No. observations	684			180		
<i>Non-small emitters</i>						
Emissions (tons of CO <sub>2</sub> )	2,318,772	7,824,052	403,134	3,084,029	8,072,944	1,282,283
No. firms	228			61		
No. observations	2,052			549		
<b>Installation-level data</b>						
<i>Coal and gas</i>						
Emissions (tons of CO <sub>2</sub> )	1,146,599	2,844,325	265,104	1,684,523	3,533,450	486,467
No. installations	462			139		
No. observations	4,158			1,251		
<i>Coal-fired installations</i>						
Emissions (tons of CO <sub>2</sub> )	3,132,774	4,795,398	1,532,082	2,010,040	3,887,297	753,381
No. installations	125			110		
No. observations	1,125			990		
<i>Gas-fired installations</i>						
Emissions (tons of CO <sub>2</sub> )	409,887	748,665	127,796	449,806	781,842	232,371
No. installations	337			29		
No. observations	3,033			261		
<b>Country-level covariates</b>						
Electricity consumption (GWh)	260,277	202,921	248,835	73,460	42,830	56,157
Renewable energy produced (GWh)	81,769	61,821	89,837	12,179	7,031	10,586
Net electricity exports (GWh)	3.00	27.45	-0.49	5.76	8.17	5.24
Real GDP (million euro)	1,300,963	1,027,274	1,103,793	209,556	133,709	159,504
<b>Cross-country covariates</b>						
Allowance price (euro/ton CO <sub>2</sub> )	8.61	3.67	7.36	8.61	3.67	7.36
Coal price (euro/MWh)	9.66	1.78	9.28	9.66	1.78	9.28
Gas price (euro/MWh)	19.51	4.74	20.03	19.51	4.74	20.03

Notes: Data on emissions are from the EUTL, data on fuel type from Open Power System Data and web search. Country-level covariates are from Eurostat, while data on allowance and fuel prices are from Bloomberg. Renewable energy produced covers production from wind, solar, biomass and hydro sources. All numbers are rounded to the last digit.

installations and then grouped by their main fuel type.<sup>16</sup> The treatment group consists

<sup>16</sup>Oil plants are not considered in the empirical analysis in this paper, as they only play a fairly minor role in the European power generation sector. Moreover, the number of oil plants in the control group is very low, leading to an insufficient sample size.

of 462 installations whose primary fuel type is either coal or gas, while there are 139 such installations in the control group. Average emissions are higher in the control group, reflecting the smaller number of gas plants.

The country-level covariates reflect differences between treatment and control group at the country level: Countries in the treatment group have larger economies on average, with larger power sectors needed to satisfy their greater demand for electricity and a higher share of consumption being satisfied by renewable energy.<sup>17</sup>

Finally, the data on fuel prices show that the fuel cost of coal-fired installations per MWh was roughly half of the cost of gas plants during our sample period, assuming that power generators in Europe approximately face the same fuel prices. This cost advantage is somewhat mitigated by the higher emission cost of coal plants, due to their greater emission intensity. However, allowance prices were low during the sample period, with a mean of less than 9 Euro per ton of CO<sub>2</sub>.

### 3.3. Matching

The regression analysis uses two main samples, the full sample as collected originally, and a matched sample. I choose a simple and transparent matching strategy, one-to-one nearest neighbor propensity score matching with replacement.<sup>18</sup> For each (sub)sample in the main analysis at the firm level, I match on the average of installations' log pre-treatment emissions, i.e. emissions during the period 2009-2012. Additionally, for the installation-level analysis installations are matched within their main fuel type, i.e. coal plants are matched with other coal plants, and analogously for gas plants.<sup>19</sup>

Figure 1 compares propensity scores for the full firm-level samples before and after matching.<sup>20</sup> There is a good overlap in the distribution of propensity scores between treated and control units in the raw samples (Figure 1(a)), ensuring an adequate extent of common support. The matched samples reflect this advantage, exhibiting closely matched propensity scores, despite the limited size of the control group (Figure 1(b)). However, part of the propensity score distribution for the treatment group does not overlap with that of the control group, for either very high or very low propensity scores. This leads to

---

<sup>17</sup>Note that renewable energy consists of wind, solar and hydro power, and of power from biofuels.

<sup>18</sup>While matching on one variable only does not necessitate propensity score matching, it is convenient for determining the area of common support.

<sup>19</sup>Unlike in the installation-level sample firms are not matched using the fuel type, as they may own both coal and gas plants.

<sup>20</sup>For comparison, Appendix A.1, Figure 8, shows the distribution of propensity scores at the installation level.

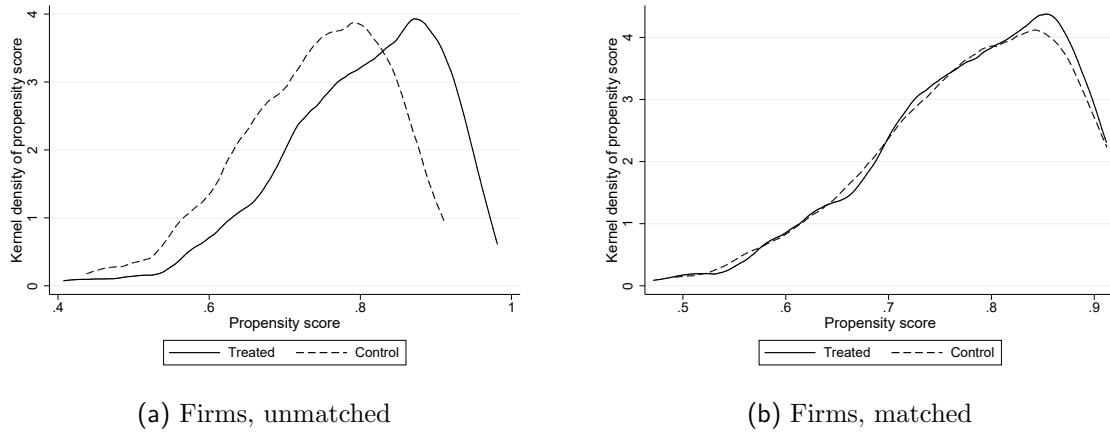


Figure 1: Propensity scores, by treatment status

Note: Distribution of propensity scores from one-to-one nearest neighbor matching with replacement, for the unmatched and matched full firm-level sample. Firms are matched on the average of log pre-treatment emissions, i.e. emissions during the period 2009-2012.

a reduction in the number of treated firms in the matched sample, compared to the full sample.

Table 2 summarizes the balance of the matching variable and of the country-level covariates used in the regression analysis in the full and matched firm-level (sub)samples. The means and variances of pre-treatment emissions differ significantly in all unmatched samples. In the matched all-firms and non-small-emitter (sub)samples both the mean and variance of pre-treatment emissions are closely aligned.<sup>21</sup> In the small-emitter subsample the match is also much-improved compared to the unmatched sample. However, the small sample size limits the ability to achieve an equally high degree of alignment as in the other two samples. The mean of the country-level covariates used in the regression analysis remain similar before and after matching, showing the need to explicitly control for these variables in the regressions.

Figure 2 shows the entire distribution of the matching variable for all three firm-level samples. In all cases, the distribution differs significantly between the treatment and control groups prior to the matching. After the matching, the distribution is closely aligned between the two groups, with quantiles close to the 45-degree axis, especially in the all-firms and non-small-emitters samples. As before, balance is more difficult to achieve in the small-emitter subsample, due to the small sample size.

<sup>21</sup>Note that the average of emissions in the treatment group changes due to the exclusion of units with propensity scores outside the common support.

Table 2: Matching summary

		Full sample	Matched sample
<i>All firms</i>			
Mean of ln(emissions 2009-2012)	Treatment	12.23	12.46
	Control	13.52	12.44
t-test equality of means (p-value)		0.00	0.89
Variance ratio		1.70	0.97
No. firms	Treatment	304	284
	Control	81	81
No. observations	Treatment	2,736	2,556
	Control	729	2,556
<i>Small emitters</i>			
Mean of ln(emissions 2009-2012)	Treatment	9.32	9.73
	Control	11.36	9.53
t-test equality of means (p-value)		0.00	0.05
Variance ratio		1.29	0.62
No. firms	Treatment	76	58
	Control	20	20
No. observations	Treatment	684	522
	Control	180	522
<i>Non-small emitters</i>			
Mean of ln(emissions 2009-2012)	Treatment	13.20	13.98
	Control	14.23	13.74
t-test equality of means (p-value)		0.00	0.96
Variance ratio		1.60	1.02
No. firms	Treatment	228	154
	Control	61	61
No. observations	Treatment	2,052	1,386
	Control	549	1,386
<i>Country-level covariates (means, all firms samples)</i>			
Electricity consumption (GWh)	Treatment	260,277	268,830
	Control	73,460	65,058
Renewable energy produced (GWh)	Treatment	81,769	83,614
	Control	12,179	11,416
Net electricity exports (GWh)	Treatment	3.00	3.00
	Control	5.76	6.60
Real GDP (million euro)	Treatment	1,300,963	1,342,970
	Control	209,556	183,132

Note: Small emitters are defined as being below the 25th percentile of the distribution of firms' 2009 emissions. Non-small emitters are defined as being at or above the 25th percentile of the distribution of firms' 2009 emissions. Means and tests for their equality across treatment and control groups are shown for the matching variable, the mean of ln(emissions) during the pre-treatment period 2009-2012. The variance ratio is defined as the ratio of sample variances of treated and control groups. The number of firms (observations) in the matched samples shows the number of firms (observations) on the common support.






Figure 2: Q-Q plots of pre-treatment emissions in treatment and control groups

Note: The graphs show quantile-quantile plots of average log emissions during 2009-2012 for all firms, small emitters and non-small emitters, by treatment status. Black dots indicate quantiles in the matched samples, dots with a white fill show quantiles in the unmatched samples. Small emitters are defined as being below the 25th percentile of the distribution of firms' 2009 emissions. Non-small emitters are defined as being at or above the 25th percentile of the distribution of firms' 2009 emissions.


### 3.4. Parallel trends

Figure 3 shows the evolution of mean emissions by treatment status, again for all three firm-level samples used in the main analysis. For better comparability across samples, the y-axes in all graphs are on the same scale for the unmatched and matched versions of each sample. The vertical line in each panel indicates the end of the pre-treatment period. Figure 3 allows us to visually assess the appropriateness of the parallel trends assumption.<sup>22</sup>


As already indicated in Table 1, the levels of average emissions differ between the treatment and control groups throughout the sample period in the unmatched samples. Emissions are always higher in the control group. The matched samples exhibit very similar pre-treatment emissions, further confirming the success of the matching strategy. Based on the visual inspection of group-level emissions the parallel trend assumption is plausible, especially for the matched samples.

Visual inspection of emissions in the post-treatment period reveals that average emissions remain closely matched for the all-firms and non-small-emitters samples, with a temporary dip in average treatment group emissions in both samples before converging back in the latter post-treatment years. The evolution of average post-treatment emissions


<sup>22</sup>Statistical tests of parallel trends are included in the regression tables.




(a) Full sample




(b) Matched sample




(c) Small emitters, unmatched



(d) Small emitters, matched



(e) Non-small emitters, unmatched



(f) Non-small emitters, matched

Figure 3: Firm level emissions, by emitter size and treatment status

Note: The y-axes show the natural logarithm of emissions by treatment status, averaged across firms in treatment and control groups. The vertical line indicates the last pre-treatment year. Small emitters are defined as being below the 25th percentile of the distribution of firms' 2009 emissions. Non-small emitters are defined as being at or above the 25th percentile of the distribution of firms' 2009 emissions.

in the control group shows no obvious break in trend, assuaging possible concerns that the modernization investment required as a condition for receiving further free allocation either did not lead to clear gains in emission efficiency, or that such gains have not yet

materialized during the sample period.

The picture differs somewhat for post-treatment emissions in the matched small-emitter subsample. After the start of the post-treatment period, average emissions in the treatment group drop below the average in the control group, until the second-to-last post-treatment year. In the last post-treatment year average emissions in the control group decline, perhaps indicating the completion of one or more modernization programs. The visual evidence suggests that a negative treatment effect in the regression analysis using the small-firms subsample should be interpreted as a lower bound, due to the drop in average emissions in the control group in the last post-treatment year.

## 4. Research design

### 4.1. Identification

The relationship of interest is the causal effect of a change in allowance allocations on emissions. However, as shown by Fowlie and Perloff (2013) for California’s RECLAIM system, a naive regression of emissions on allocations produces spurious results due to the endogeneity of installation-level allocations. Both current allocations and current emissions are likely to be correlated with unobservable historical emissions. Current emissions are correlated with historical emissions due to path dependency induced by the technology of each installation, while the correlation of current allocations and historical emissions stems from the allocation rule in the EU ETS, which determined allocations based on historical pre-ETS emissions.

To identify the causal effect, I exploit variation in the allocation rule induced by a policy change that became effective in 2013, in the middle of the sample period. While electricity producers in most EU member states lost the bulk of their free allocation after the end of EU ETS Phase II in 2012, significant free allocation – although also decreased compared to previous years – continued for power producers in eight member states under the 10c rule. Installations having mostly lost free allocation constitute the treatment group, while those under the 10c rule are the control group. In the matched sample average firm-level allocations in the treatment group fell from more than 1.5 million allowances in 2012 to about 100 thousand allowances in 2013, a decrease by more than 93% (Figure 4(b)).<sup>23</sup> Free allocation also diminished in the control group, as allocation levels did not remain

---

<sup>23</sup>The remaining share of free allocation is explained by the fact that producers in the treatment group continue receiving some free allocation for the production of heat.

at pre-2013 levels under the 10c rule. However, the drop in allocation levels in 2013, by about 44%, was far less pronounced. Allocations also remain far higher in the remainder of the sample period.<sup>24</sup>

In both the full and matched samples, the treatment group experiences a similar evolution of free allocation. For the control group, the average allocation level is higher in the full sample than in the matched one, although the relative drop from 2012 to 2013 is very similar, with a decrease in allocated allowances by about 46% in the full sample (Figure 4(a)).

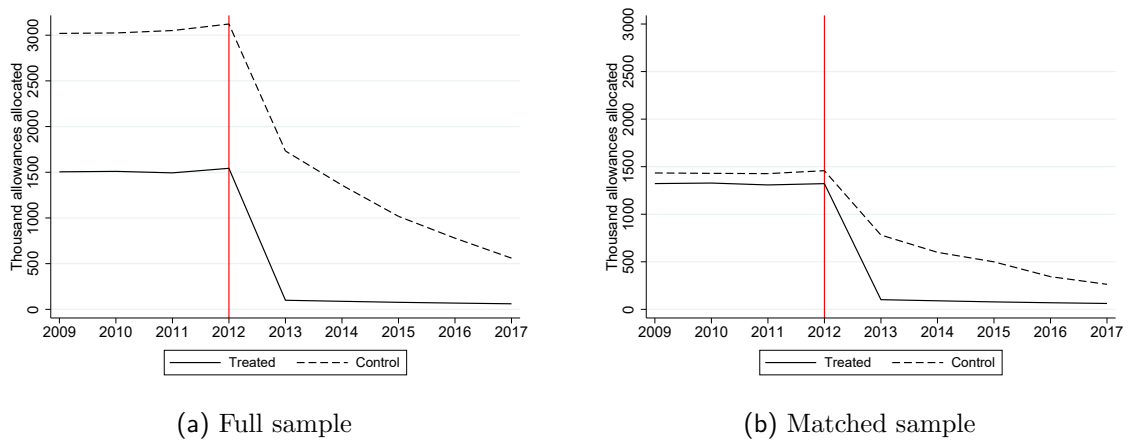


Figure 4: Firm-level allowance allocations, by treatment status

Note: The y-axes show the number of allowances allocated, by treatment status. Y-axes are on the same scale for better comparability.

The divergence in allocation levels and therefore in producers’ demand for additional allowances leads to a difference in total compliance costs across treatment and control groups. If transaction costs depend on the amount of allowances purchased, treated producers bear a relatively greater burden of transaction costs starting in 2013. If this is the case, treated producers will purchase fewer allowances than socially optimal and instead abate more than socially optimal. Their emissions are expected to decrease compared to power producers in the control group.

#### 4.2. Validity of the control group

For the purpose of this study it is important to understand whether free allocation under the 10c rule may have led to deviations in emission patterns in the control group compared to its emissions under free allocation based on grandfathering, as was practiced until 2012.

<sup>24</sup>Allocation patterns are comparable at the installation level, see Appendix A.2, Figure 9.

In case of deviations it is important to understand the direction in which emissions of firms from the control group would be affected to determine the direction of a possible bias in the regression analysis.

The stated main purpose of free allocation under the 10c rule was to support firm's modernization expenditures. Even though the selection and thus additionality of modernization projects is not entirely clear due to a lack of transparency, intuition suggests that post-2012 free allocation which came with strings attached should lead to some gains in emission intensity for facilities in the control group, compared to the pre-2012 regime in which free allocation came with no strings attached. I.e., emissions in the control group may be expected to be lower in the post-2012 period compared to what they would have been under grandfathering. It is unclear how large any such difference would be and how quickly it would manifest. According to IEA (2019) investment projects in thermal power generation on average take about 4 years until completion, suggesting that any gains in emission intensity are likely to be realized several years into the post-treatment period. The descriptive evidence on control group emissions also does not suggest a break in the emission pattern in the control group (see Section 3.4), with the exception of the small emitters subsample, in which a clear drop in emissions is visible in 2017, five years after the change in the allocation regime and in line with the information in IEA (2019). Nevertheless, a cautious assumption is that emissions in the control group decreased somewhat due to the 10c allocation rule, compared to what they would have been under grandfathering. If treated generators behave according to the transaction cost hypothesis, i.e. if they reduce their emissions due to the treatment, lower emissions in the control group would somewhat underestimate the true effect of moving from free allocation to full auctioning in the treatment group. I.e., the empirical analysis estimates a lower bound on the true treatment effect.

Another potential concern is that emitters in the control group may have a greater emission intensity on average, either for strategic reasons or because power producers in the control group may be less sophisticated than those in the treatment group. Higher emissions in the control group *per se* would not lead to bias in the DID analysis, which focuses on within-producer *changes* in emissions. Only a change in average emission intensity in response to the 10c program, perhaps due to greater scrutiny by market observers, could be a potential concern. However, if anything, emissions in the control group can be expected to decline due to potential added scrutiny under the 10c program.

This would, again, introduce a bias against the direction of the expected treatment effect, again leading to a lower bound interpretation of the treatment effect.

Overall, as potential biases due to emission decisions of producers in the control group under the 10c rule are against the expected direction of the treatment effect, power producers in the control group can be viewed as a valid counterfactual for treated producers. Additionally, due to the long lead times for investment projects and the opaque nature of the modernization conditionality for free allocation under the 10c rule I expect any potential bias to be small.

A further potential concern are spillover effects between treatment and control groups. In principle, power producers in the treatment group could be at a competitive disadvantage compared to producers in the control group due to the allowance subsidy to the latter. Producers in the control group could use its cost advantage to export electricity to the treatment group. Such an increase in electricity trade due to the allocation policy change would be a threat to identifying the treatment effect, as emissions in the control group would bias the regression analysis in the same direction as the expected treatment effect.




Figure 5: Net electricity exports from control to treatment groups

Source: Eurostat.

Examining the evolution of electricity trade between the treatment and control groups indicates whether such a spillover may have occurred. Note that options to trade electricity are limited by the need for available transmission infrastructure, limiting potential spillover

effects. For this reason the spillover issue is reduced to trade with the countries bordering the treatment group: Poland, the Czech Republic, Slovakia and Hungary on the exporter side, with Germany and Austria as possible destination countries.

Figure 5 shows net exports (i.e. exports net of imports) of electricity from the above control group countries to the treatment group countries. We observe positive net electricity exports from the control group to the treatment group during the pre-treatment period. During the post-treatment period net exports diminish and control group countries become net importers of electricity from the treatment group, in line with rising power production from renewables in the treatment group, especially in Germany. This assuages concerns about spillovers due to the change in allocation policy, but also highlights the need to control for renewable energy production and cross-border electricity trade in the regression analysis.

### 4.3. Empirical model

I estimate the following DID equation as the preferred specification:

$$y_{it} = \alpha_i + \lambda_t + \rho D_{it} + X'_{ct}\beta + \epsilon_{it} \quad (1)$$

$y_{it}$  is the natural logarithm of emissions of producer  $i$  in year  $t$ .  $D_{it}$  is a dummy variable indicating when treatment is switched on for producer  $i$ . Accordingly,  $\rho$  is the average treatment effect on the treated (ATT).  $\alpha_i$  are producer fixed effects and  $\lambda_t$  are year fixed effects.  $X_{ct}$  contains the country-level covariates capturing the relevant outcomes on the electricity market of country  $c$  in which the fossil-fueled installation  $i$  is located: electricity consumption,<sup>25</sup> the quantity of renewable energy produced, and net power exports. Additionally, I control for macroeconomic activity by including GDP. Standard errors are clustered at the firm level to robustify against serial correlation in emissions (Bertrand et al., 2004). For the installation-level analysis, this additionally allows for correlation of emission decisions across installations within the same firm.

I also run the following alternative model:

---

<sup>25</sup>I thus assume that demand for electricity is inelastic.

$$y_{it} = \alpha_i + \nu Post_t + \rho D_{it} + X'_{ct}\beta + Z'_t\gamma + \epsilon_{it} \quad (2)$$

Equation (2) replaces year fixed effects in Equation (1) with matrix  $Z_t$ , which contains allowance prices and the prices of coal and gas, which are assumed to be common to all countries in our sample and thus only vary in the time series dimension. As year fixed effects are excluded in Equation (2), the  $Post$  indicator is not absorbed by them and thus remains in the estimation.

I run the regression analysis on different samples: The main analysis uses the full firm-level sample (Section 5.1) and two firm-level subsamples, small emitters and non-small emitters (Section 5.2).<sup>26</sup> Further analysis uses installation level data (Sections 6.1) and differentiates further by fuel type and owner type (Section 6.2).

Additional analysis in Section 5.3 considers the dynamics of the treatment effect by estimating the following equation at the firm level:

$$y_{it} = \alpha_i + \lambda_t + \sum_{j=-m}^q \rho_j D_{it,t=0+j} + X'_{ct}\beta + \epsilon_{it} \quad (3)$$

Equation 3 is similar to Equation 1, except that the treatment effect indicator is disaggregated into  $m$  leads (pre-treatment periods in this terminology) and  $q$  lags (post-treatment periods), with the last pre-treatment period left out as the baseline period.  $t = 0$  indicates the start of the post-treatment period. In our case, there are four pre-treatment periods (2009-2012) and five post-treatment periods (2013-2017). This set-up allows for a formal test of the parallel trends assumption: If the parallel trends assumption holds,  $\rho_j = 0 \forall j < 0$ , i.e. the pre-treatment policy indicators will be jointly not different from zero. The p-values of this test are included for each regression in the results tables in Section 5. Moreover, using coefficient estimates on the post-treatment coefficients of the policy variable we can evaluate the evolution of the treatment effect over time.

---

<sup>26</sup>See Section 3.2 for a definition of the two subsamples.



## 5. Main results

### 5.1. Effect of change in allocation on firm-level emissions

Table 3 presents the baseline results for the firm-level sample.<sup>27</sup> Columns (1)-(3) show the estimates for the full sample of all available firms and columns (4)-(6) present the corresponding results for the matched sample. Columns (1) and (4) contain a simple specification only including the policy indicator, firm fixed effects and year fixed effects. Specifications (2) and (5) add the following country-level controls: Amount of energy produced from renewable sources, electricity consumption, and net electricity exports. Specifications (3) and (6) add allowance and fuel prices as cross-country controls. They do not include year fixed effects, which would absorb the cross-country controls. These specifications have the advantage of explicitly modelling the role of input prices. However, as they do not include year fixed effects, they do not control for other sources of cross-country shocks to firms' emissions. Specifications (2) and (5) are the most demanding ones, as they control not only for allowance and fuel prices but for any cross-country factors. They are therefore the preferred specifications. F-tests of joint significance of the pre-treatment coefficients ("leads") in Equation (3) (e.g. Autor, 2003), provide a test of the validity of the common trend assumption.

In the full sample, switching from free allocation to full auctioning is estimated to significantly decrease installations' emissions in all three specifications, by between 9.5% and 14.4%.<sup>28</sup> In the matched sample point estimates are of the same sign as in the full sample. However, they are far smaller, while standard errors are larger. The change in allocation is estimated to decrease emissions by between 5.6% and 8%. Neither of the point estimates is statistically significant.

The results for the electricity market covariates included in the regressions are largely as expected. A greater consumption of electricity is related to higher firm-level emissions, as are larger net exports of electricity. In the full sample an increase in power generation from renewables is related to a decrease in emissions, which is also expected, as greater production from low-marginal cost renewables drives out fossil-based suppliers. This effect is insignificant in the matched sample, reflecting the larger share of firms from the control

---

<sup>27</sup>Recall that that the definition of firms in this paper follows the definition used by the European Commission in the EUTL, i.e. each firm is limited to one country only.

<sup>28</sup>Note that the R-squared is very high in all regressions. This is due to the inclusion of firm fixed effects. Doing so controls for, in particular, generation technology, which explains a large share of installations' emissions.

Table 3: Baseline firm-level results

	(1)	(2)	(3)	(4)	(5)	(6)
	Full sample			Matched sample		
ATT	-0.095** (0.040)	-0.144*** (0.045)	-0.140*** (0.044)	-0.066 (0.059)	-0.056 (0.066)	-0.080 (0.057)
ln(RE)		-0.135* (0.074)	-0.195*** (0.066)		0.124 (0.125)	0.039 (0.101)
ln(Electricity consumption)		0.988** (0.425)	1.740*** (0.390)		1.390*** (0.540)	1.681*** (0.543)
Net electricity exports		0.003*** (0.001)	0.004*** (0.001)		0.002 (0.001)	0.002** (0.001)
ln(GDP)		-0.965*** (0.320)	-1.400*** (0.320)		-1.042*** (0.324)	-1.497*** (0.346)
EUA price			0.001 (0.003)			0.007** (0.004)
Coal price			0.012** (0.005)			-0.004 (0.007)
Gas price			-0.004* (0.002)			-0.002 (0.003)
Installation FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	N	Y	Y	N
F-test leads (p-value)	0.03	0.10	0.02	0.40	0.22	0.77
Obs	3,465	3,465	3,465	5,112	5,112	5,112
No. firms	385	385	385	352	352	352
R <sup>2</sup>	0.984	0.985	0.984	0.984	0.984	0.984

Notes: OLS regressions of the natural log of firm-level emissions on a policy dummy switching from 0 to 1 after the start of the treatment period for treated installations, in 2013. Specifications (1) and (3) include the policy dummy, installation fixed effects and year fixed effects. Specifications (2) and (5) additionally control for country-level electricity market outcomes and GDP. Specifications (3) and (6) add allowance prices and fuel prices, while not controlling for year fixed effects. All regressions include F-tests of joint significance of switching on the policy variable in each pre-treatment year. Robust standard errors clustered at the firm level are in parentheses. \*\*\*, \*\* and \* indicate significance at the 1%, 5% and 10% level, respectively.

group in the matched sample, which operate in power markets with a far lower penetration of renewable energy. An increase in GDP is related to a drop in emissions, perhaps due to an ability to devote more resources to climate policy when the overall economy does well. A higher allowance price is either not estimated to affect emissions or is even positively correlated with emissions, in line with the previous finding in the literature that firms can fully pass through their emission cost (Fabra and Reguant, 2014).

Based on the results for the matched sample I fail to reject the independence property for power producers under the EU ETS. These results suggest that transaction costs are not substantial enough to cause a dependence between allocations and emissions. However, independence is rejected for the full sample. This contradictory result motivates differentiating between small emitters and non-small emitters in the further analysis, as small emitters are more likely to be resource constrained compared to the average firm,

making them more vulnerable to transaction costs.

An additional insight from the estimations in Table 3 is that the specifications including global controls and excluding year fixed effects yield results that are fairly comparable to the preferred specification. Controlling for allowance and fuel prices appears to cover a relevant share of unobserved cross-country heterogeneity. However, given the generic advantage of the preferred specification and its better performance with respect to parallel pre-treatment trends, it remains the focus of the analysis in the following sections.

## 5.2. Heterogeneous effects by emitter size

In this section I split the firm-level sample and compare the effect of changing the allocation level for the subsample of small emitters with the subsample of non-small emitters.<sup>29</sup> The sample split is motivated by the insight that larger emitters may be better able to absorb transaction costs, e.g. additional staffing resources required to participate in allowance auctioning. If transaction costs are significant, emissions of non-small emitters should react less strongly than those of small emitters, as the former can reap economies of scale while the latter are constrained in doing so. If transaction costs are very low, neither small nor non-small emitters should exhibit a dependence between allocations and emissions. Table 4 contains the results for the matched sample and follows Table 3 with respect to specifications.<sup>30</sup>

I find that small emitters significantly decrease their emissions due to the loss in allocation. The preferred specification yields a statistically significant decrease in emissions by about 19%. The result for the matched sample of small emitters is in agreement with the analogous result for the unmatched sample (Appendix A.3, Table 7), where small emitters are estimated to decrease their emissions by about 23% due to the loss of free allocation.

In contrast, no statistically significant effect of the loss in allocation is found for larger emitters in the matched sample. Point estimates are also much smaller than for small emitters, in the range of -4.2% to -5.9%. Note that while a significant decrease in emissions is found for larger emitters in the full sample, the effect is much smaller in magnitude than for small emitters (Appendix A.3, Table 7).

These results are in line with the hypothesis that small emitters are resource-constrained and therefore less able to fully absorb the additional cost of full auctioning for allowances

---

<sup>29</sup>See Section 3.2 for a description of the construction of the small and non-small emitters subsamples.

<sup>30</sup>For comparison, Table 7 (Appendix A.3) shows the results for the full sample.

Table 4: Small and non-small emitters, matched sample

	(1)	(2)	(3)	(4)	(5)	(6)
	Small emitters			Non-small emitters		
ATT	-0.156 (0.103)	-0.192** (0.097)	-0.206*** (0.072)	-0.052 (0.042)	-0.042 (0.042)	-0.059 (0.041)
ln(RE)		0.218 (0.319)	0.253 (0.278)		0.008 (0.080)	-0.071 (0.069)
ln(Electricity consumption)		0.759 (1.671)	1.377 (1.694)		1.170** (0.521)	1.362*** (0.478)
Net electricity exports		-0.002 (0.003)	-0.001 (0.003)		0.001 (0.001)	0.002 (0.001)
ln(GDP)		-2.451** (1.048)	-3.475*** (0.883)		-0.599** (0.299)	-0.818*** (0.287)
EUA price			0.015 (0.009)			0.004 (0.003)
Coal price			-0.015 (0.026)			0.001 (0.005)
Gas price			-0.003 (0.010)			-0.0004 (0.002)
Installation FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	N	Y	Y	N
F-test leads (p-value)	0.46	0.98	0.01	0.69	0.84	0.88
Obs	1,044	1,044	1,044	2,772	2,772	2,772
No. firms	62	62	62	203	203	203
R <sup>2</sup>	0.829	0.833	0.821	0.966	0.967	0.966

Notes: OLS regressions of the natural log of firm-level emissions on a policy dummy switching from 0 to 1 after the start of the treatment period for treated installations, in 2013. Samples are divided into small and non-small emitters. Small emitters are defined as being below the 25th percentile of the distribution of firms' 2009 emissions. Non-small emitters are defined as being at or above the 25th percentile of the distribution of firms' 2009 emissions. Specifications (1) and (3) include the policy dummy, installation fixed effects and year fixed effects. Specifications (2) and (5) additionally control for country-level electricity market outcomes and GDP. Specifications (3) and (6) add allowance prices and fuel prices, while not controlling for year fixed effects. All regressions include F-tests of joint significance of switching on the policy variable in each pre-treatment year. For better readability, heating and cooling degree days are rescaled by a factor of 1,000. Robust standard errors clustered at the firm level are in parentheses. \*\*\*, \*\* and \*\* indicate significance at the 1%, 5% and 10% level, respectively.

without distortions to their emission decisions. Such distortions are not detected for larger firms, who appear to be able to better internalize these transaction costs. However, note that small emitters as defined in this paper represent a total of 3.2 million tons of CO<sub>2</sub> emissions over the entire sample period 2009-2017, out of a total of 7.7 trillion tons for the power sector under the EU ETS during the same period, i.e. about 0.4% of total power sector emissions under the EU ETS. This suggests that although a significant number of firms exhibit a dependence between allocations and emissions, the overall allowance market is unlikely to be affected. This intuition is borne out by the empirical results for non-small emitters and for the sector-level results using the pooled sample (Section 5.1),

both of which suggest that the independence property holds for the power sector as a whole.

The result for small emitters adds a new dimension to the existing literature on the independence property in cap-and-trade schemes, which either considers the entire allowance market (Fowlie and Perloff, 2013) or (to my understanding) larger emitters in the Spanish power market Reguant and Ellerman (2008). At the same time, my analysis confirms the main results from the existing literature that finds overall support for the independence property.

As a policy implication, my result for small emitters validates thinking further about ways of easing the burden for small emitters, an issue recognized in current EU ETS rules, which foresees an opt-out for small emitters under certain conditions (EU, 2018). The issue of how to further decrease costs for market participants is likely to remain relevant in the policy debate, as further sectors with more small and medium-sized emitters are likely to be added to the EU ETS in the future.

### 5.3. Dynamics

This section analyzes the dynamics of the treatment effect in the matched firm-level samples. Treatment effect dynamics for the analogous cases in the full sample are broadly in line with those for the matched sample (Appendix A.4, Figure 10). Figures 6 and 7 depict the point estimates and 95% confidence bands of the annually disaggregated treatment effect as in Equation (3). I consider the preferred specifications for the three cases presented in Sections 5.1 and 5.2, i.e. the sample of all firms, the small emitter subsample and the non-small emitter subsample.

Statistically significant anticipation effects are absent in all cases, as the leads (i.e. pre-treatment periods) of the policy indicator are both individually and jointly indistinguishable from zero. With respect to the post-treatment periods, the analysis of the pooled sample (Figure 6) reveals a short-term statistically significant decrease in emissions in the second post-treatment year. Point estimates revert back to zero thereafter and are statistically insignificant in all other post-treatment years.

Results deviate when splitting the sample according to emitter size. For small emitters, (Figure 7(a)), emissions decline in every post-treatment period, except for a very imprecisely estimated increase in the final post-treatment period. The latter coincides with a decrease in emissions in the control group in 2017 also visible in the raw data (Figure 3(d)).




Figure 6: Treatment effect dynamics, all firms and matched sample

Note: OLS regression of the natural log of firm-level emissions on a policy dummy switching from 0 to 1 in every pre and post-treatment year, country-level electricity market outcomes,  $\ln(\text{GDP})$ , installation fixed effects and year fixed effects. Treatment effects and 95% confidence intervals are depicted on the y-axis, using robust standard errors clustered at firm level to construct confidence intervals. The vertical line in each panel indicates 2012, the end of the pre-treatment period. 2012 is the baseline year. Number of observations: 5,112; number of firms: 352.

Treatment effect dynamics for small emitters show a significant decrease in emissions due to the loss in free allocation, consistent with the transaction cost hypothesis. Results differ markedly in the non-small emitter subsample (Figure 7(b)). Point estimates are close to zero and not statistically significant in any post-treatment period. The emissions of the group of firms representing the vast majority of aggregate power sector emissions (>99%) emissions are found to be statistically invariant to the change in their allocation levels, in line with the independence property.

The result for non-small emitters is therefore consistent with the interpretation that transaction costs under the EU ETS are small enough to be absorbed by the majority of firms representing most of the sector's emissions, so that the independence property holds for the – from a climate perspective – relevant share of the power sector under the EU ETS. However, for 25% of firms the independence property fails, even though this failure only affects a very small share of both the power sector's emissions and the overall




Figure 7: Treatment effect dynamics by emitter size, matched samples

Note: OLS regressions of the natural log of firm-level emissions on a policy dummy switching from 0 to 1 in every pre and post-treatment year, country-level electricity market outcomes,  $\ln(\text{GDP})$ , installation fixed effects and year fixed effects. Treatment effects and 95% confidence intervals are depicted on the y-axis, using robust standard errors clustered at firm level to construct confidence intervals. The vertical line in each panel indicates 2012, the end of the pre-treatment period. 2012 is the baseline year. Number of observations small emitter subsample: 1,044; number of firms small emitter subsample: 62. Number of observations non-small emitter subsample: 2,772; number of firms non-small emitter subsample: 203.

allowance market.

## 6. Further results

### 6.1. Baseline installation-level results

In this section and in Section 6.2 I further probe the firm-level results by exploiting installation-level emission data, information about installations' main fuel type and whether an installation belongs to a firm owning only one or multiple installations regulated under the EU ETS. Table 5 contains the installation-level analogue of the results in Table 3.<sup>31</sup>

In the full sample, switching from free allocation to full auctioning is estimated to significantly decrease installations' emissions by between 7.9% and 11.8%, respectively. The point estimates for the full sample at the installation level are smaller in magnitude than those at the firm level, but they do indicate a possible dependence between allocations and emissions, as in the full firm-level sample. In the matched sample point estimates are close to zero and never statistically significant.

<sup>31</sup>To economize on table size and manuscript length the installation-level results tables focus on the treatment effect. Full results are available upon request.

Table 5: Baseline installation-level results

	(1)	(2)	(3)	(4)	(5)	(6)
	Full sample			Matched sample		
ATT	-0.079** (0.035)	-0.107*** (0.040)	-0.118*** (0.039)	0.042 (0.083)	0.013 (0.103)	0.004 (0.095)
Country-level controls	N	Y	Y	N	Y	Y
Global controls	N	N	Y	N	N	Y
Installation FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	N	Y	Y	N
F-test leads (p-value)	0.21	0.40	0.23	0.95	0.49	0.81
Obs	5,409	5,409	5,409	8,154	8,154	8,154
No. firms	601	601	601	544	544	544
R <sup>2</sup>	0.975	0.975	0.975	0.973	0.973	0.973

Notes: OLS regressions of the natural log of installation-level emissions on a policy dummy switching from 0 to 1 after the start of the treatment period for treated installations, in 2013. Specifications (1) and (3) include the policy dummy, installation fixed effects and year fixed effects. Specifications (2) and (5) additionally control for country-level electricity market outcomes and GDP. Specifications (3) and (6) add allowance prices and fuel prices, while not controlling for year fixed effects. All regressions include F-tests of joint significance of switching on the policy variable in each pre-treatment year. Robust standard errors clustered at the firm level are in parentheses. \*\*\*, \*\* and \* indicate significance at the 1%, 5% and 10% level, respectively.

The results for the full sample suggest that the independence property also holds at the installation level, not just at the firm level. However, as in the firm-level sample, the results for the full sample indicate that some installations' emission decisions may not be independent of their allocation levels.

## 6.2. Installation-level results by fuel type and firm structure

The results reported in Table 6 approximate the firm-level distinction between small and large emitters at the installation level, with the restriction that the sample size becomes too small to match installations simultaneously on fuel type and on whether the parent company is a small emitter. Due to this constraint on the installation-level analysis I approximate small emitters by differentiating between installations owned by firms with only one installation under the EU ETS, as opposed to installations belonging to firms with more than one such installation. This approximation is imperfect, as there are cases of one-installation firms with fairly high emissions, e.g. one-installation firms owning a coal plant.

The results are somewhat similar to the firm-level results for small emitters: installations from one-installation firms are estimated to significantly decrease emissions in response to the loss in allocation, by about 13.7% (Table 6, column (2)). In contrast, the



Table 6: Installation-level results, single-installation and multi-installation firms

	(1)	(2)	(3)	(4)	(5)	(6)
	Full sample			Matched sample		
	All	Single	Multi	All	Single	Multi
ATT	-0.107*** (0.040)	-0.137** (0.054)	-0.054 (0.060)	0.013 (0.102)	-0.002 (0.131)	0.102 (0.099)
Country-level controls	Y	Y	Y	Y	Y	Y
Global controls	N	N	N	N	N	N
Installation FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
F-test leads (p-value)	0.40	0.05	0.06	0.49	0.09	0.04
Obs	5,409	2,700	2,709	8,154	3,996	3,114
No. installations	601	300	301	544	257	228
R <sup>2</sup>	0.975	0.979	0.964	0.973	0.974	0.972

Notes: Notes: OLS regressions of the natural log of installation-level emissions on a policy dummy switching from 0 to 1 after the start of the treatment period for treated installations, in 2013. Specifications (1) and (3) include the policy dummy, installation fixed effects and year fixed effects. Specifications (2) and (5) additionally control for country-level electricity market outcomes and GDP. Specifications (3) and (6) add allowance prices and fuel prices, while not controlling for year fixed effects. All regressions include F-tests of joint significance of switching on the policy variable in each pre-treatment year. Robust standard errors clustered at the firm level are in parentheses. \*\*\*, \*\* and \* indicate significance at the 1%, 5% and 10% level, respectively.

point estimate for installations from multiple-installation firms is much smaller in magnitude and statistically insignificant. In the matched sample, I do not find a significant dependence between the allocation level and emissions for this group of installations.

As a further robustness check, Appendix A.5, Table 8 further differentiates the estimates provided in Table 6 for each fuel type. Point estimates are always smaller for installations from single-installation firms than for their multiple-installation analogues. However, neither subsample reveals a significant dependence between allocations and emissions, likely due to a lack of statistical power given the small sample sizes after the double sample split.

## 7. Summary and conclusions

This paper provides an empirical test of the independence property postulated under the Coase Theorem. Specifically, I analyze whether greenhouse gas emissions are independent from allocations for power producers under the EU Emissions Trading System (EU ETS). This paper is – to my knowledge – the first to conduct an analysis of the independence property across firms and installations in multiple countries of a large-scale cap-and-trade

program targeting greenhouse emissions. To my knowledge, this is also the first paper to differentiate emitters by size while paying special attention to emission decisions by small emitters.

The empirical analysis pursues a difference-in-differences strategy. The causal effect of allowance allocations on emissions is identified by leveraging a change in allocation policy. Due to a policy shift power producers lost most of their freely allocated allowances in the majority of EU member states (treatment group), while producers in a number of European countries retained substantial free allocation under a special provision – the 10c rule – (control group).

The main regression analysis uses a panel of European power producing firms from 2009 to 2017 and controls for the main electricity market outcomes and macroeconomic activity, as well as for cross-country drivers of emissions such as allowance and fuel prices. In the main analysis I pay particular attention to small emitters, firms that due to resource constraints may be less able to absorb transaction costs without facing distorted emission decisions. I run estimations on full and matched samples and also estimate treatment effect dynamics. Additional regressions using disaggregated installation-level data, for coal and gas-fired installations, further elaborate on the main firm-level analysis.

I fail to reject independence between allocations and emissions for the power sector overall and for the 75% of firms representing more than 99% of power sector emissions. However, I do find that the independence property fails for the lowest quartile of emitters. The loss in free allocation is estimated to decrease emissions of small emitters by about 20%. The analysis of treatment effect dynamics suggests that the effect on their emissions remains throughout the sample period. These results suggest that transaction costs in the EU ETS can be absorbed by most firms such that allowance allocations and emissions remain independent. Although a subset of firms – small emitters – exhibits distorted emission decisions, these distortions do not significantly affect the independence property at the sector level.

The analysis in this paper has important policy implications, e.g. for adding new sectors to the EU ETS or for policy makers considering the introduction of a new cap-and-trade scheme using the EU ETS as a blueprint. My results suggest that – as long as the setting is comparable to the one studied in this paper – policy makers may use free allocation in the political bargaining process without significantly distorting emissions decisions under the cap-and-trade program. However, my finding that the independence

property fails for small emitters signals a need for further efforts to decrease the costs of participation in the EU ETS or in schemes modelled after the EU ETS. Doing so would likely also increase the political acceptance of cap-and-trade. Moreover, when deciding in favor of free allocation, the allocation system should be designed to minimize incentives of firms to distort the program in order to secure higher free allocation in the future. Finally, policy makers should consider questions of environmental justice when pursuing free allocation.

One limitation of this paper is that a treatment effect analysis as performed here is only feasible for the power sector. While I conclude that the power sector – the most important sector under the EU ETS – does not contribute to a failure of the independence property in the EU ETS at large, independence may be compromised due to emitters in manufacturing sectors covered by the same program. Indeed, a dependence between allocation and emissions is expected if free allocation to manufacturing firms fulfils its intended purpose of protecting against carbon leakage as opposed to mostly providing firms with windfall profits. Future research must determine whether the independence property in the EU ETS fails due to manufacturing firms.

## References

- Abrell J., Kosch M., and Rausch S. (2019) How Effective Was the UK Carbon Tax?- A Machine Learning Approach to Policy Evaluation. *CER-ETH-Center of Economic Research at ETH Zurich Working Paper*, 19: 317.
- Ambec S. and Ehlers L. (2016) Regulation via the Polluter-pays Principle. *The Economic Journal*, 126(593): 884–906.
- Autor D. H. (2003) Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing. *Journal of Labor Economics*, 21(1): 1–42.
- Bayer P. and Aklin M. (2020) The European Union emissions trading system reduced CO2 emissions despite low prices. *Proceedings of the National Academy of Sciences*, 117(16): 8804–8812.
- Bertrand M., Duflo E., and Mullainathan S. (2004) How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics*, 119(1): 249–275.

- Boehringer C. and Lange A. (2005) On the Design of Optimal Grandfathering Schemes for Emission Allowances. *European Economic Review*, 49(8): 2041–2055.
- Calel R. (2020) Adopt or Innovate: Understanding Technological Responses to Cap-and-Trade. *American Economic Journal: Economic Policy*, 12(3): 170–201.
- Calel R. and Dechezlepretre A. (2016) Environmental policy and directed technological change: evidence from the European carbon market. *Review of Economics and Statistics*, 98(1): 173–191.
- Carlson C., Burtraw D., Cropper M., and Palmer K. L. (2000) Sulfur dioxide control by electric utilities: What are the gains from trade? *Journal of Political Economy*, 108(6): 1292–1326.
- CMW (2016) Fossil Fuel Subsidies from Europe’s Carbon Market. *Carbon Market Watch Policy Briefing April 2016*.
- Coase R. H. (1960) The Problem of Social Cost. *The Journal of Law and Economics*, 3: 1–44.
- Colmer J., Martin R., Muûls M., Wagner U. J. et al. (2020) Does Pricing Carbon Mitigate Climate Change? Firm-Level Evidence from the European Union Emissions Trading Scheme. *Firm-Level Evidence from the European Union Emissions Trading Scheme (November 5, 2020)*.
- Dechezleprêtre A., Gennaioli C., Martin R., Muûls M., and Stoerk T. (2019) Searching for carbon leaks in multinational companies.
- Eden A., De Oca M. M., and Haug C. (2019) Striving to Keep ETS Simple. *International Carbon Action Partnership (ICAP)*.
- Ellerman A. D., Marcantonini C., and Zaklan A. (2016) The European Union emissions trading system: ten years and counting. *Review of Environmental Economics and Policy*, 10(1): 89–107.
- EU (2001) Directive 2001/80/EC of the European Parliament and of the Council of 23 October 2001 on the limitation of emissions of certain pollutants into the air from large combustion plants (Large Combustion Plant Directive). *On line at: <http://eur-lex.europa.eu/legalcontent/EN/TXT/PDF>*.


- (2009) Directive 2009/29/EC OF The European Parliament and of the Council of 23 April 2009 Amending Directive 2003/87/EC so as to Improve and Extend the Greenhouse Gas Emission Allowance Trading Scheme of the Community. *Official Journal of the European Union*, L 140/63.
- (2011) Communication from the Commission - Guidance document on the optional application of Article 10c of Directive 2003/87/EC (2011/C 99/03). *Official Journal of the European Union*, C 9/9.
- (2018) Directive 2018/29/EC OF The European Parliament and of the Council of 14 March 2018 Amending Directive 2003/87/EC to enhance cost-effective emission reductions and low-carbon investments, and Decision (EU) 2015/1814. *Official Journal of the European Union*, L 76/3.
- Fabra N. and Reguant M. (2014) Pass-through of emissions costs in electricity markets. *American Economic Review*, 104(9): 2872–99.
- Fleurbaey M. and Maniquet F. (2018) Optimal income taxation theory and principles of fairness. *Journal of Economic Literature*, 56(3): 1029–79.
- Fowlie M. and Perloff J. M. (2013) Distributing pollution rights in cap-and-trade programs: are outcomes independent of allocation? *Review of Economics and Statistics*, 95(5): 1640–1652.
- Hahn R. W. (1984) Market power and transferable property rights. *The Quarterly Journal of Economics*, 99(4): 753–765.
- Hanemann W. M. (1991) Willingness to pay and willingness to accept: how much can they differ? *The American Economic Review*, 81(3): 635–647.
- Harstad B. and Eskeland G. S. (2010) Trading for the Future: Signaling in Permit Markets. *Journal of Public Economics*, 94(9-10): 749–760.
- Heindl P. (2012) Transaction costs and tradable permits empirical evidence from the EU emissions trading scheme. *ZEW Discussion Papers*, 12.
- (2017) The Impact of Administrative Transaction Costs in the EU Emissions Trading System. *Climate Policy*, 17(3): 314–329.

- Hepburn C. J., Quah J. K.-H., and Ritz R. A. (2013) Emissions trading with profit-neutral permit allocations. *Journal of Public Economics*, 98: 85–99.
- IEA (2019) World Energy Investment 2019. *International Energy Agency (IEA), Paris*, <https://www.iea.org/reports/world-energy-investment-2019>.
- Jaraitė J., Convery F., and Di Maria C. (2010) Transaction costs for firms in the EU ETS: lessons from Ireland. *Climate Policy*, 10(2): 190–215.
- Jaraitė J., Jong T., Kažukauskas A., Zaklan A., and Zeitlberger A. (2016) Ownership links and enhanced EUTL dataset. Available at <https://cadmus.eui.eu/handle/1814/64596>.
- Kahneman D., Knetsch J. L., and Thaler R. H. (1990) Experimental tests of the endowment effect and the Coase theorem. *Journal of Political Economy*, 98(6): 1325–1348.
- Liski M. and Montero J.-P. (2010) Market power in an exhaustible resource market: The case of storable pollution permits. *The Economic Journal*, 121(551): 116–144.
- Löfgren Å., Burtraw D., Wråke M., and Malinovskaya A. (2018) Distribution of Emissions Allowances and the Use of Auction Revenues in the European Union Emissions Trading System. *Review of Environmental Economics and Policy*, 12(2): 284–303.
- Martin R., Muuls M., de Preux L., and Wagner U. (2014) Industry Compensation under Relocation Risk: A Firm-Level Analysis of the EU Emission Trading Scheme. *American Economic Review*, 104(8): 2482–2508.
- Meunier G., Montero J.-P., and Ponsard J.-P. (2018) Output-Based Allocations in Pollution Markets with Uncertainty and Self-Selection. *Journal of Environmental Economics and Management*, 92: 832–851.
- Montero J.-P., Sanchez J. M., and Katz R. (2002) A market-based environmental policy experiment in Chile. *The Journal of Law and Economics*, 45(1): 267–287.
- Montgomery W. D. (1972) Markets in licenses and efficient pollution control programs. *Journal of Economic Theory*, 5(3): 395–418.
- Naegele H. and Zaklan A. (2019) Does the EU ETS cause carbon leakage in European manufacturing? *Journal of Environmental Economics and Management*, 93: 125–147.
- Petrick S. and Wagner U. (2014) The Impact of Carbon Trading on Industry: Evidence from German Manufacturing Firms. *Kiel Working Paper*, 1912: .

- Popa D. and Cepraga I. (2012) The Article 10c Application by Romania: Comments Regarding the Application of Romania for Transitional Allocation Free of Charge to Installations for Electricity Production. *Bankwatch Network/Greenpeace*.
- Reguant M. and Ellerman A. D. (2008) Grandfathering and the endowment effect: An Assessment in the context of the Spanish National Allocation Plan. *Center for Energy and Environmental Policy Research, Cambridge, Massachusetts*.
- Rosendahl K. E. (2008) Incentives and Prices in an Emissions Trading Scheme with Updating. *Journal of Environmental Economics and Management*, 56(1): 69–82.
- Stavins R. N. (1995) Transaction costs and tradeable permits. *Journal of Environmental Economics and Management*, 29(2): 133–148.
- UNFCCC (2015) Adoption of the Paris agreement. *Proposal by the President (Draft Decision), FCCC/CP/2015/L.9/Rev.1*.

## A. Appendix

### A.1. Distribution of installation-level propensity scores



(a) Installations, unmatched

(b) Installations, matched

Figure 8: Propensity scores, by treatment status

Note: Distribution of propensity scores from one-to-one nearest neighbor matching with replacement, for the unmatched and matched full installation-level sample. Installations are matched on the average of log pre-treatment emissions, i.e. emissions during the period 2009-2012, within the main fuel type (coal and gas).



## A.2. Installation-level allowance allocations, by treatment status




Figure 9: Installation-level allowance allocations, by treatment status

Note: The y-axes show the number of allowances allocated, by treatment status. Y-axes are on the same scale for better comparability.

### A.3. Heterogeneity in firm-level emissions by size, full sample

Table 7: Small and non-small emitters, full sample

	(1)	(2)	(3)	(4)	(5)	(6)
	Small emitters			Large emitters		
ATT	-0.196** (0.098)	-0.226** (0.110)	-0.206* (0.105)	-0.061 (0.042)	-0.106** (0.046)	-0.106** (0.046)
ln(RE)		-0.123* (0.160)	-0.134 (0.140)		-0.122 (0.077)	-0.190*** (0.067)
ln(Electricity consumption)		1.959** (0.950)	3.315*** (0.915)		0.556 (0.463)	1.157*** (0.401)
Net electricity exports		0.002 (0.003)	0.002 (0.003)		0.002** (0.001)	0.003** (0.001)
ln(GDP)		-2.380* (1.322)	-3.369*** (1.010)		-0.575** (0.278)	-0.885*** (0.274)
EUA price			0.0004 (0.008)			0.0004 (0.003)
Coal price			0.026* (0.015)			0.010** (0.005)
Gas price			-0.008 (0.006)			-0.003 (0.002)
Installation FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	N	Y	Y	N
F-test leads (p-value)	0.01	0.04	0.001	0.50	0.37	0.37
Obs	864	864	864	2,601	2,601	2,601
No. firms	96	96	96	289	289	289
R <sup>2</sup>	0.942	0.943	0.942	0.976	0.976	0.976

Notes: OLS regressions of the natural log of firm-level emissions on a policy dummy switching from 0 to 1 after the start of the treatment period for treated installations, in 2013. Samples are divided into small and non-small emitters. Small emitters are defined as being below the 25th percentile of the distribution of firms' 2009 emissions. Non-small emitters are defined as being at or above the 25th percentile of the distribution of firms' 2009 emissions. Specifications (1) and (3) include the policy dummy, installation fixed effects and year fixed effects. Specifications (2) and (5) additionally control for country-level electricity market outcomes and GDP. Specifications (3) and (6) add allowance prices and fuel prices, while not controlling for year fixed effects. All regressions include F-tests of joint significance of switching on the policy variable in each pre-treatment year. Robust standard errors clustered at the firm level are in parentheses. \*\*\*, \*\* and \* indicate significance at the 1%, 5% and 10% level, respectively.

#### A.4. Treatment effect dynamics by firm size




Figure 10: Treatment effect dynamics, full sample

Note: OLS regressions of the natural log of firm-level emissions on a policy dummy switching from 0 to 1 in every pre and post-treatment year, country-level electricity market outcomes,  $\ln(\text{GDP})$ , installation fixed effects and year fixed effects. Treatment effects and 95% confidence intervals are depicted on the y-axis, using robust standard errors clustered at firm level to construct confidence intervals. The vertical line in each panel indicates 2012, the end of the pre-treatment period. 2012 is the baseline year. Number of observations full sample: 3,465; number of firms full sample: 385. Number of observations small emitter subsample: 864; number of firms small emitter subsample: 96. Number of observations non-small emitter subsample: 2,601; number of firms non-small emitter subsample: 289.

## A.5. Installation-level results by fuel type and parent firm type

Table 8: Installation-level results by fuel type, single-installation vs. multi-installation firms

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Full sample				Matched sample			
	Single coal	Single gas	Multi coal	Multi gas	Single coal	Single gas	Multi coal	Multi gas
ATT	-0.021 (0.054)	-0.087 (0.125)	0.112 (0.082)	-0.051 (0.148)	-0.006 (0.070)	0.008 (0.160)	0.092 (0.100)	0.059 (0.121)
Country-level controls	Y	Y	Y	Y	Y	Y	Y	Y
Global controls	N	N	N	N	N	N	N	N
Installation FE	Y	Y	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
F-test leads (p-value)	0.47	0.02	0.01	0.55	0.94	0.11	0.01	0.19
Obs	720	1,980	1,395	1,314	648	3,348	1,602	1,512
No. installations	80	220	155	146	57	200	130	98
R <sup>2</sup>	0.983	0.983	0.963	0.947	0.978	0.966	0.960	0.934

Notes: OLS regressions of the natural log of installation-level emissions on a policy dummy switching from 0 to 1 after the start of the treatment period for treated installations, in 2013. Specifications (1) and (3) include the policy dummy, installation fixed effects and year fixed effects. Specifications (2) and (5) additionally control for country-level electricity market outcomes and GDP. Specifications (3) and (6) add allowance prices and fuel prices, while not controlling for year fixed effects. All regressions include F-tests of joint significance of switching on the policy variable in each pre-treatment year. Robust standard errors clustered at the firm level are in parentheses. \*\*\*, \*\* and \* indicate significance at the 1%, 5% and 10% level, respectively.