

1850

Discussion
Papers

Coase and Cap-and-Trade: Evidence on the Independence Property from the European Electricity Sector

Aleksandar Zaklan

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

IMPRESSUM

© DIW Berlin, 2020

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 Electricity Sector

Aleksandar Zaklan*

First version: February 2016

This version: March 2020

This paper provides an empirical test of the Coase Theorem. I analyze whether emissions are independent from allowance allocations in the electricity sector regulated under the EU's Emissions Trading System (EU ETS). Exogenous variation in levels of free allocation for power producing installations enables a difference-in-differences strategy. The analysis reveals that a change in allocation levels does not significantly affect emissions, either at the plant or firm level. However, I identify an adjustment period with temporarily lower emissions for some firms. The results suggest that policy makers may use free allocation as a political tool without compromising the program's cost-effectiveness.

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 Petrick, Brendan Restrepo, Alejandro 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).¹ 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, making cost-effectiveness robust to arbitrary distributions of emission 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).²

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's 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 simplifies the calculation for policy makers. They do not need to be concerned about distorting the scheme through free allocation, and 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 their cost-effectiveness may be compromised, in which case a liberal use of free allocation should be avoided to limit distortions.

This paper tests whether emissions are independent from allowance allocations for electricity-producing installations regulated under the world's largest multilateral cap-

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

²Independence may also be affected by imperfect competition in the allowance market (Hahn, 1984; Liski and Montero, 2010) or behavioral disparities between agents' willingness to pay and the willingness to accept (Kahneman et al., 1990; Hanemann, 1991).

and-trade scheme, the EU’s Emissions Trading System (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.³ Dependence between allocations and emissions for this group of producers would indicate a loss in the program’s cost-effectiveness. 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 a panel of gas-fired and coal-fired installations – the producers of the bulk of emissions in the European power sector – from 2009 to 2017.⁴ The causal effect of allowance allocations on installations’ emissions is identified by exploiting exogenous variation in installation-level allowance allocations: 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 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 2013,⁵ a special provision – the so-called 10c rule – preserved a significant amount of free allocation for installations in eight EU member states (EU, 2011).⁶ The analysis thus compares emissions decisions of installations who lost most of their free allocation (treatment group) with those who continue receiving substantial free allocation under the 10c rule (control group).⁷

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

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

⁴Oil-fired power generation is excluded from this analysis as it, first, does not play a large role in Europe and, second, an analysis is not feasible due to the small number of oil-fired installations in the control group.

⁵Electricity generators continue receiving a limited number of allowances for other activities, especially the generation of heat (see Section 2).

⁶These countries are Bulgaria, Cyprus, Czech Republic, Estonia, Hungary, Lithuania, Poland, and Romania. Free allocation under the 10c rule will continue until at least 2030 (EU, 2015).

⁷UK installations are excluded from the analysis, as the UK the 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).

purchased, the divergence in allocation levels will expose treated installations to a differential in transaction costs. If treated installations 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.

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 include installation fixed effects.⁸ Furthermore, I control for 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, either implicitly through year fixed effects, or explicitly, by including allowance and fuel prices. Estimations are run on two different samples, raw and matched. I match on the average of pre-treatment emissions, within generation technology.

In addition to investigating the independence property using the pooled sample of all coal and gas plants, a full assessment also requires the analysis of important heterogeneities, to test whether a possible dependence between allocations and emissions may be masked in the full sample. The analysis considers heterogeneity in the following dimensions: First, I analyze the role of parent companies' exposure to the EU ETS by separating out the subsample of installations belonging to firms with multiple installations under the EU ETS. This dimension is important, as more exposed companies may be better able to spread (fixed) transaction costs over several installations, so that the additional cost may be lower for installations belonging to larger players. Companies with deeper experience trading under the EU ETS may also be less prone to behavioral biases. Second, I aggregate the installation-level data to the firm level, to better understand whether some firms, especially smaller emitters, may exhibit dependence between allocations and emissions, which may not be apparent in the disaggregated analysis at the installation level.

Based on the main analysis I fail to reject the independence property, i.e. I do not find that a change in the level of free allocation causes a change in emissions. This finding holds for the preferred specification in practically all subsamples: at the installation and firm level in all matched samples, and in most cases for the raw samples. One exception is the raw firm-level sample including single-installation firms. There, I find a significant negative relationship between allowance allocations and emissions of treated firms. This

⁸Including installation fixed effects nests country fixed effects, as installations cannot relocate.

would be in line with the prediction from theory that the presence of transaction costs causes installations not receiving enough free allocation to cover all their emissions to abate more than would be cost-effective in a transaction-cost-free setting. To better understand the source of this result I conduct additional analysis of treatment effect dynamics, which reveals a short-term decrease in emissions in the firm-level sample. The effect is transitory, and point estimates of the treatment effect converge back towards zero. Based on this additional analysis I conclude that transaction costs in the EU ETS are fairly low, and mostly fixed rather than variable. Otherwise, we would observe a more permanent dependence between allocations and emissions.⁹ Some firms may take some time to adapt to the new environment of lower allocation levels, e.g. by hiring additional staff to optimize firms' allowance auctioning strategy or secondary market purchasing. However, it appears that transaction costs – or other factors affecting the behavior of emitters – are not substantial enough to significantly divert equilibrium emissions of the power sector as a whole from their cost-effective levels.

The findings in this paper are consistent with the conclusion that well-administered cap-and-trade schemes may exist at large scale without incurring significant losses in cost-effectiveness. This result is of policy relevance, 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 sacrificing the cost-effectiveness of the cap-and-trade program, as long as transaction costs are comparable to those facing power producers under the EU ETS. Of course, even though cost-effectiveness losses may be avoided free allocation will have distributional implications.

To the best of my knowledge, this paper is the first to comprehensively analyze the independence property in a large-scale cap-and-trade program for greenhouse gas emissions using a quasi-experimental approach. In terms of empirical strategy, it is most similar to Fowlie and Perloff (2013), who consider the independence property for nitrogen oxide (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 NO_x and sulfur oxides (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

⁹This finding is in line with the literature on transaction costs in the EU ETS, which concludes that these mostly consist of fixed costs and are quantitatively moderate (Jaraitė et al., 2010).

abatement option.¹⁰ 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, and by expanding the analysis beyond Spanish coal plants and beyond the EU ETS pilot phase.¹¹ A further relevant but distinct contribution is Carlson et al. (2000), who estimate abatement cost curves for the U.S. SO₂ trading program and evaluate the extent to which cost economies were reaped during its early years. Montero et al. (2002) show that transaction costs can be high if an allowance market does not have a strong institutional foundation.

This paper is also related to the literature analyzing the effectiveness of cap-and-trade programs, and especially the effectiveness of the EU ETS. Petrick and Wagner (2014) estimate the effectiveness of the EU ETS with respect to abatement of German manufacturing firms, while Dechezleprêtre et al. (2019) and Naegele and Zaklan (2019) are examples of contributions evaluating its effectiveness with regard to international emission abatement, i.e. whether abatement in Europe is negated by carbon leakage. Cael and Dechezlepretre (2016) 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 in addition to 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 cost.

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. Section 5 contains the results, Section 6 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 starting in 2013 power producing installations shall lose their free allocation and be subject to full

¹⁰Carbon 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.

¹¹The latter point is relevant, as the EU ETS changed significantly since its pilot phase. See, e.g., Ellerman et al. (2016) for a description of the main changes.

auctioning of emissions allowances. Cogeneration units may continue to receive some free allocation, but only for the generation of heating and cooling. District heating plants also continue receiving free allocation for the emissions involved in generating heat. This means that the electricity sector transitioned rapidly from widespread free allocation to mostly full auctioning between the 2012 and 2013 compliance years.

Article 10c lays out the rules for continued free allocation to installations generating electricity during Phase III, while Communication 2011/C 99/03 (EU, 2011) fleshes out how Article 10c is implemented. Member states' governments may 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. Power plants either in existence or *physically initiated* by December 31, 2008 were eligible to receive allowances for free.¹² The stated goal of continued free allocation under Article 10c is to support installations in the affected states with modernization measures and to avoid undue increases in electricity prices as a result of a shift to full auctioning. Allowances allocated for free under the 10c rule were subtracted from each member states auctioning contingent, i.e. member states opting for free allocation under the 10c rule incurred the consequence of decreased revenue from auctioning allowances.

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. From the perspective of the regression analysis, this opt-out underscores that self-selection takes place at the country level. Each member state using the provision under the 10c rule had to submit a list of possible investment options in modernizing its power sector 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 have to report to the Commission on the progress of investments each year. The lists of investment options submitted by the affected member states and approved by the Commission exceeded the

¹²The meaning of physically initiated is not entirely clear, but does not imply that the plant must have been under construction by December 31, 2008. Plants that would come online years later would also be eligible for free allocation under the 10c rule.

value of free allocation, so that installations' operators had several options to conduct their modernization measures and were therefore not tied to specific measures.

For the purpose of this study it is of interest whether free allocation under the 10c rule was substantially different from the way in which free allocation was conducted previously or whether emission patterns could be expected to be affected by modernization projects. While the process under the 10c rule was not very transparent, evidence from the grey literature suggests that investment projects under the 10c rule mainly subsidized the continued production of fossil-fueled power and thus did not lead to systematic decrease in emissions of the affected installations (CMW, 2016). The proposed lists of investments were criticized for supporting fossil fuel generation facilities by incumbent companies, and that a large portion of proposed projects did not pursue the goal of decreasing emissions (Popa and Cepraga, 2012).

The evidence on the implementation of the 10c rule suggests that the allocation process under the 10c rule was substantially similar to the allocation based on historical emissions prior to 2013. It is therefore appropriate to assume that 10c installations' emissions trajectories are a fair representation of the treated installations' emissions paths, had they not lost most of their free allowances after 2012.

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.¹³ The sample period is 2009-2017.¹⁴

Note the UK Carbon Price Support (CPS), which introduced a minimum price for the

¹³Focusing 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.

¹⁴Observations 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.

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.

Data on the main fuel type of installations is obtained from the Open Power System Data platform,¹⁵ supplemented by a manual online search. The country-level electricity market covariates electricity consumption, the amount of electricity produced from renewable energies, heating degree days and cooling degree days are from Eurostat. So is data on GDP per capita. Data on allowance and fuel prices are from Bloomberg.

3.2 Descriptive overview

Table 1 provides a descriptive overview of the dataset. I present the installation-level data grouped by their main fuel type, either coal or gas.¹⁶ Information on the main fuel type and installed capacity is available for 842 out of 862 installations identified as belonging to the power sector using the sector mapping provided by the European Commission (Jaraitė et al., 2016). The treatment group consists of the majority of EU member states, where installations lost the vast majority of their free allocation starting in 2013 (see Section 4.1). The control group is substantially smaller and consists of installations that continue receiving free allocation – although generally less than up to 2012 – after 2013.

The subsample of coal plants consists of 96 treated and 102 control installations. Emissions are larger on average in the treatment group. The distributions of emissions are skewed toward a small number of large emitters. This is why emissions enter the regression analysis in logarithmic form. The subsample of gas plants consists 361 installations in the treatment group and 30 in the control group. However, despite the limited size of the control group, mainly for gas plants, the distributions of pre-treatment emissions and capacity among treated and control gas plants exhibit a good overlap in terms of propensity scores (see Section 3.3).

With respect to the firm-level sample, I follow the European Commission’s definition of a firm, denoted as ”account holders” in the EUTL data. Mean emissions in the treatment group are lower in the firm-level sample than in the sub-sample of coal plants. This reflects

¹⁵<https://open-power-system-data.org/>

¹⁶Oil 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.

Table 1: Descriptive overview

	Treatment group			Control group		
	Mean	Std. Dev.	Median	Mean	Std. Dev.	Median
<i>Coal-fired installations:</i>						
Emissions (tons of CO ₂)	3,260,348	5,105,004	1,526,284	2,012,822	3,982,209	724,223
No. installations	96			102		
No. observations	864			918		
<i>Gas-fired installations:</i>						
Emissions (tons of CO ₂)	465,838	1,240,244	128,783	601,722	1,187,828	199,996
No. installations	361			30		
No. observations	3,249			270		
<i>Firm-level data:</i>						
Emissions (tons of CO ₂)	1,620,070	6,850,825	175,718	2,827,335	7,397,630	626,227
No. firms	297			79		
No. observations	2,673			711		
<i>Country-level covariates:</i>						
Electricity consumption (GWh)	253,448	192,859	241,869	81,366	44,379	56,385
Renewable energy produced (GWh)	82,088	58,556	89,246	13,343	7,146	10,791
Heating degree days	3,012	1,254	2,909	3,219	388	3,286
Cooling degree days	88	114	17	54	57	28
Real GDP per capita (euro)	31,519	8,474	33,280	11,281	3,429	10,510
<i>Cross-country covariates:</i>						
Allowance price (Euro per ton CO ₂)	8.61	3.66	7.36	8.61	3.66	7.36
Coal price (Euro per MWh)	9.66	1.78	9.28	9.66	1.78	9.28
Gas price (Euro per 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.

the fairly large number of single-installation firms operating gas plants.

The country-level covariates reflect differences between treatment and control group at the country level: Countries in the treatment group are richer 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.¹⁷ Treated countries are also somewhat warmer on average, with a smaller number of heating degree days and a larger number of cooling degree days.

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

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

emission intensity. However, allowance prices were low during the sample period, with an average of less than 9 Euro per ton of CO₂.

3.3 Matching

The regression analysis uses two samples, the raw sample as collected originally, and a matched sample. I choose a simple and transparent matching strategy, one-to-one propensity-score matching with replacement. I match on the average of installations' log pre-treatment emissions, i.e. emissions during the period 2009-2012. Installations are matched within their main fuel type, i.e. coal plants are matched with other coal plants, and analogously for gas plants. The matching produces a sample with two groups of similar installations, i.e., two groups of coal plants of similar 2012 emissions, and likewise for gas-fired installations.




Figure 1: Propensity scores, by treatment status

The firm-level sample is also matched using the average of pre-treatment emissions. However, unlike in the installation-level sample firms are not matched using the fuel type, as they may own both coal and gas plants.

Figure 1 compares propensity scores before and after matching, for both the installation-level and firm-level samples. There is a good overlap in the distribution of propensity scores between treated and control units in the raw samples (Figure 1, panels (a) and (c)), 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, panels (b) and (d)).

Table 2: Matching summary

		Treatment	Control	t-test (pvalue)
<i>Coal-fired installations:</i>				
Mean of ln(emissions 2009-2012)	Unmatched	14.21	13.58	0.00
	Matched	14.27	14.26	0.99
Main fuel type		<i>Exact match</i>		
No. installations on common support		95	102	
No. observations on common support		855	918	
<i>Gas-fired installations:</i>				
Mean of ln(emissions 2009-2012)	Unmatched	11.69	12.13	0.24
	Matched	11.81	11.79	0.98
Main fuel type		<i>Exact match</i>		
No. installations on common support		351	30	
No. observations on common support		3,159	270	
<i>Firms:</i>				
Mean of ln(emissions 2009-2012)	Unmatched	12.10	13.49	0.00
	Matched	12.70	12.70	0.99
No. firms on common support		249	79	
No. observations on common support		2,241	711	

Note: 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. Additionally, installations are matched within each main fuel type.

Table 2 summarizes the balance of the matching variable – the average of log pre-treatment emissions – in the raw and matched samples. Levels of pre-treatment emissions differ in the raw installation-level sample for coal plants, while being more similar in the unmatched gas plant sample. In the matched installation-level sample there are no significant differences between treatment and control groups with respect to pre-treatment emissions of either coal or gas plants. In the raw firm-level sample the average of pre-treatment emissions differs significantly between treatment and control groups, while being very similar in the matched sample.

3.4 Parallel trends

Figure 2 shows the evolution of mean emissions by treatment status, for the installation-level and firm-level samples. The graphs compare emissions in the raw and matched samples. For better comparability across samples, the y-axes in all graphs are on the same scale. The vertical line in each panel indicates the end of the pre-treatment period. Figure 2 allows us to visually assess the appropriateness of the parallel trends assumption.¹⁸




Figure 2: Installation and firm level emissions, by treatment status

Note: The y-axis of all graphs shows the natural logarithm of emissions by treatment status, averaged across installations in treatment and control groups. For better comparability y-axes are on the same scale. The vertical line indicates the last pre-treatment year.

The levels of average emissions differ between the treatment and control groups throughout the sample period in the raw samples. Emissions are higher in the control group in both the installation-level and firm-level samples. 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,

¹⁸Formal tests of the parallel trends assumption are included in the regression tables.

especially for the matched samples.

Visual inspection of emissions in the post-treatment period reveals that emissions remain closely matched in the installation-level sample. In the firm-level sample emissions exhibit a decline in the treatment group compared to the control group, before converging back in the latter post-treatment years.

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 at the country level 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 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 installation-level allocations in the treatment group fell from more than 940,000 allowances in 2012 to about 64,000 allowances in 2013, a decrease by more than 93% (Figure 3(a)).¹⁹. In contrast, free allocation also diminished in the control group, as allocation levels did not remain at pre-2013 levels under the 10c rule. However, the drop in allocation levels in 2013, by about 40%, was far less pronounced. Allocations also remain far higher in the remainder of the sample period and installations in the control group will continue

¹⁹Note that some installations in the treatment group continue receiving free allocation for the production of heat.

receiving significant free allocation for the foreseeable future. Free allocation under the 10c rule will continue until at least 2030 (EU, 2015).

In both the raw 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 raw sample than in the matched one, although the relative drop from 2012 to 2013 is similar, with a decrease in allocated allowances by about 44% (Figure 3(b)).




Figure 3: 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.

The divergence in allocation levels and therefore in installations' 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 installations bear a relatively greater burden of transaction costs starting in 2013. If this is the case, treated installations will purchase fewer allowances than socially optimal and instead abate more than socially optimal. Their emissions are expected to decrease compared to installations in the control group.

4.2 Empirical model

I estimate the following DID equation as the preferred specification:

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

y_{ict} is the natural logarithm of emissions of installation i in country c and year t . D_{ct} is a dummy variable indicating when treatment is switched on in country c . Accordingly,

ρ is the average treatment effect on the treated. α_i are installation fixed effects and λ_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,²⁰ the quantity of renewable energy produced, heating degree days and cooling degree days. Additionally, I control for macroeconomic activity by including GDP per capita. Heating and cooling degree days are included to account for within-year variation in electricity demand. Standard errors are clustered at the firm level to robustify against serial correlation in emissions (Bertrand et al., 2004) and to allow for correlation of emission decisions across installations within the same firm.

I also run the following alternative model:

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

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 main regression analysis on different samples: The main analysis uses the full sample of all coal and gas fired installations (Section 5.1). I then consider the subsample of installations from multiple-installation firms to explore the role of exposure of installations' parent companies to the EU ETS (Section 5.2). Doing so is relevant, as more exposed companies may be better able to spread (fixed) transaction costs over several installations. Companies with deeper experience trading in the EU ETS may also exhibit less of a behavioral bias. Finally, I aggregate the installation-level data to the firm level (Section 5.3). This helps me understand whether a possible dependence between allocations and emissions by some firms, especially smaller emitters, may be masked in the disaggregated analysis at the installation level.

Additional analysis in Section 5.4 considers the dynamics of the treatment effect by estimating the following equation:

²⁰I thus assume that demand for electricity is inelastic.

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

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.

5 Results

5.1 Full sample

Table 3 presents the results for the full installation-level sample. Columns (1)-(3) show the estimates for the matched sample and columns (4)-(6) present the corresponding results for the raw sample. Columns (1) and (4) contain a simple specification only including the treatment effect indicator, installation fixed effects and year fixed effects. Specifications (2) and (5) add country-level controls, and 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 modeling the role of input prices. However, as they cannot include year fixed effects, they do not control for other sources of cross-country shocks to installations' emissions. The most demanding specifications, in columns (3) and (6), include country-level controls, installation fixed effects and year fixed effects and thus control not just for allowance and fuel prices but for all 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), confirm the validity of the common trend assumption for the preferred specification.

In the raw sample, switching from free allocation to full auctioning is estimated to

Table 3: Installation-level results, full sample

	(1)	(2)	(3)	(4)	(5)	(6)
	Matched sample			Raw sample		
ATT	-0.012 (0.078)	-0.018 (0.080)	0.038 (0.088)	-0.098 (0.036)	-0.101 (0.042)	-0.064 (0.043)
ln(RE)		0.146 (0.136)	0.370 (0.178)		-0.190 (0.059)	-0.029 (0.071)
ln(Electricity consumption)		0.487 (0.735)	0.526 (0.717)		0.753 (0.381)	0.651 (0.381)
Heating degree days		0.197 (0.060)	0.188 (0.090)		0.181 (0.027)	0.183 (0.054)
Cooling degree days		0.048 (0.308)	1.245 (0.497)		0.155 (0.161)	0.897 (0.211)
ln(GDP per capita)		-0.463 (0.545)	-0.046 (0.528)		-0.172 (0.346)	0.035 (0.365)
EUA price		0.018 (0.005)			0.007 (0.003)	
Coal price		-0.010 (0.008)			-0.001 (0.005)	
Gas price		-0.0001 (0.004)			-0.0003 (0.002)	
Installation FE	Y	Y	Y	Y	Y	Y
Year FE	Y	N	Y	Y	N	Y
F-test leads (p-value)	0.72	0.36	0.24	0.28	0.06	0.84
Obs	8,028	8,028	8,028	5,301	5,301	5,301
R ²	0.974	0.974	0.975	0.977	0.977	0.977

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 (2) include the policy dummy, installation fixed effects and year fixed effects. Specifications (2) and (4) additionally control for country-level electricity market outcomes, allowance prices and fuel prices, without controlling for year fixed effects. Specifications (3) and (5) control for country-level electricity market outcomes and 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.

significantly decrease installations' emissions in the less demanding specifications (Table 3, columns (4) and (5)), by 9.8% and 10.1%, respectively. The point estimate drops to -6.4% in the preferred specification and becomes statistically insignificant.²¹ In the matched sample point estimates are far smaller. The change in allocation is estimated to decrease emissions by 1.2% and 1.8% in the less demanding specifications (Table 3, columns (1) and (2)). These point estimates are economically and statistically insignificant. In the preferred specification, the point estimate is 3.8%, again not significant statistically.

²¹Note that the R-squared is very high in all regressions. This is due to the inclusion of installation fixed effects. Doing so controls for, in particular, generation technology, which explains a large share of installations' emissions.

Based on these results I fail to reject the independence property for power producers under the EU ETS. The results suggest that transaction costs, or behavioral biases by decision-makers at electricity-producing installations under the EU ETS, are not substantial enough to cause a loss in cost-effectiveness.

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 by not entirely comparable to the preferred specification. Controlling for allowance and fuel prices while disregarding unobserved sources of cross-country shocks covers a relevant portion of unobserved cross-country heterogeneity, but does not account for it entirely. Given the advantages of the preferred specification the analysis focuses on it in the following sections.

5.2 The role of parent companies' exposure to the EU ETS

This section compares the treatment effect in the full sample with the effect for the subsample of installations whose parent firms are more strongly exposed to the EU ETS, in the sense that more than one installation affected by the EU ETS belongs to the same firm. Multiple-installation firms may be able to spread fixed transaction costs, e.g. additional staffing resources required to participate in allowance auctioning, over several installations. If transaction costs under the EU ETS consist of mostly fixed costs and if those are significant, emissions of installations belonging to multiple-installation firms should react less than those in the full sample, as their additional cost burden would be relatively smaller. If transaction costs are low, the results for the full installation-level sample and those excluding plants from single-installation firms should be comparable.

Table 4 contains the results. Columns (1) and (2) of Table 4 reproduce the main results for the preferred specification from Table 3, while columns (3) and (4) contain the results for the subsample of installations belonging to multiple-installation firms. While eliminating installations from single-installation firms halves the number of observations in the matched sample and decreases it by about 47% in the raw sample, it eliminates a smaller share of total emissions from the sample – about 27% – as installations belonging to single-installation firms tend to be smaller emitters.

In the matched sample the point estimate of the treatment effect for the group of installations belonging to multiple-installation firms is very similar to the estimate in the full sample, 4.3% vs. 3.8%. In both cases the estimate is statistically insignificant. In the raw sample, the treatment effect is estimated to be closer to zero in the multiple-

Table 4: Installation-level results, full sample vs. installations from multiple-installation firms

	(1)	(2)	(3)	(4)
	Full sample		Multi-installation	
	Matched	Raw	Matched	Raw
ATT	0.038 (0.088)	-0.064 (0.043)	0.043 (0.099)	-0.012 (0.062)
Country-level controls	Y	Y	Y	Y
Global controls	N	N	N	N
Installation FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
F-test leads (p-value)	0.24	0.84	0.48	0.20
Obs	8,028	5,301	3,978	2,808
R ²	0.975	0.977	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. All columns contain the preferred specification, corresponding to specifications (3) and (6) in Table 3. 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.

installation sample, at -1.2% instead of -6.4%. In both cases the effect is statistically insignificant.

Removing installations belonging to single-installation firms from the sample shows that the remaining plants – those belonging to multiple-installation firms – yield similar parameter estimates and are also statistically insignificant. These results are also in line with the independence property holding. Looking at the entire fleet of installations or restricting our attention to firms with deeper exposure to the EU ETS does not make a difference with respect to the independence property. This suggests that that transaction costs facing power producers under the EU ETS are fairly low.

5.3 Firm-level results

I now consider the firm level. The firm-level analysis is motivated by the possibility that a possible dependence of emissions on allowance allocations may be masked in the installation-level analysis. It may be that plants belonging to firms owning many installations are run more efficiently, i.e. in line with independence property. If this is the case the installation-level analysis would give more weight to such firms than aggregating installations to the firm level and running the analysis at that level.

Having aggregated emission data to the firm level, I run the analysis for two samples:

a full sample including all firms, and a subsample of larger players only, i.e. only those firms owning multiple installations under the EU ETS (Table 5). Again, I consider a raw and a matched sample.

Table 5: Firm-level results, full sample vs. installations from multiple-installation firms

	(1)	(2)	(3)	(4)
	Full sample		Multi-installation	
	Matched	Raw	Matched	Raw
ATT	-0.074 (0.056)	-0.116 (0.048)	-0.065 (0.047)	-0.063 (0.063)
Country-level controls	Y	Y	Y	Y
Global controls	N	N	N	N
Installation FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
F-test leads (p-value)	0.28	0.50	0.76	0.47
Obs	4,482	3,384	1,098	729
R ²	0.978	0.984	0.988	0.987

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. All columns contain the preferred specification, corresponding to specifications (3) and (6) in Table 3. 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.

In the full raw sample I estimate that the decrease in allocation level causes a – statistically significant – decrease in emissions by 11.6%, violating the independence property (Table 5, column (2)). This result would be consistent with an interpretation that transaction costs affect firms’ emission decisions. It would be in line with the prediction from theory that firms with a level of free allocation insufficient to fully cover their emissions will abate more than is socially optimal.

The point estimate of the treatment effect is insignificant in all other cases, i.e. when using the matched version of the full firm-level sample or when removing single-installation firms in the raw sample. In the matched full sample the point estimate is -7.4% and insignificant. The estimate for multiple-installation firms is very similar for the matched and raw samples, with a point estimate of -6.5% and -6.3%, respectively, both being statistically insignificant.

The firm-level results broadly confirm the results at the installation level. However, the result for the raw full firm-level sample raises concerns that some firms may exhibit a dependence between their allocation level and emissions. The result is consistent with the interpretation that firms may decrease their emissions by abating more in response to the

additional transaction cost from participating in allowance auctioning or due to increased purchases on the secondary market. In Section 5.4 I evaluate the dynamics of the treatment effect to better understand the source of the result for the full firm-level sample.

5.4 Dynamics

This section analyzes the dynamics of the treatment effect. Figure 4 depicts the point estimates and 95% confidence bands of the annually disaggregated treatment effect as in Equation (3). I consider four different cases using the matched version of the sample: The full installation-level sample (Figure 4(a)), the subsample of installations belonging to multiple-installation firms (Figure 4(b)), the full firm-level sample (Figure 4(c)) and the subsample of multiple-installation firms (Figure 4(d)).




Figure 4: Treatment effect dynamics

Note: Treatment effects and 95% confidence intervals are depicted on the y-axis, using robust standard errors clustered at firm level to construct confidence intervals. For better comparability y-axes are on the same scale. The vertical line in each panel indicates the end of the pre-treatment period.

Statistically significant anticipation effects are absent. In all cases, the leads (i.e. pre-

treatment periods) of the policy indicator are jointly indistinguishable from zero. With respect to the post-treatment periods, the results deviate between the installation level and the firm level. At the installation level, the treatment effect is close to zero and statistically insignificant in every post-treatment year, both in the full sample (Figure 4(a)) and in the sample with installations from multiple-installation firms only (Figure 4(b)). At the firm level, I obtain negative and significant point estimates shortly after the loss in free allocation, in 2014, which revert back to zero thereafter. Qualitatively, this result is similar for the full firm-level sample (Figure 4(c)) and the subsample of larger emitters only (Figure 4(d)).²² As the temporary decrease in emissions is more pronounced in the full sample, the point estimate is larger in the full samples in the main analysis, and significant in the raw full firm-level sample (Table 5). However, when looking at the post-treatment period as a whole and when considering the matched sample of comparable firms the emissions of the power sector are found to be independent of their allocation levels, as shown in the main regression analysis.

The analysis of treatment effect dynamics confirms our intuition from the main firm-level analysis. It appears that it took some firms, especially smaller emitters, a period of time to adjust to the loss of their free allocation, e.g. because they had to hire additional staff. These results are in line with the findings in the literature that transaction costs in the EU ETS mostly consist of fixed costs and are moderate (Jaraitė et al., 2010; Naegele, 2018). The results on treatment effect dynamics suggest that variable transaction costs do not play an important role for power producers under the EU ETS. Significant variable transaction costs would have led to a more permanent – instead of a temporary – divergence in emissions between treated and control units. Overall, the additional analysis of dynamics broadly confirms our main result that emissions are independent of allocations for power producers, while identifying the existence of an adjustment period for some firms, during which they temporarily lowered their emissions.

6 Summary and conclusions

This paper provides an empirical test of the Coase Theorem. Specifically, I analyze whether greenhouse gas emissions are independent from allocations for power producers under the EU's Emissions Trading System (EU ETS). This paper is – to my knowledge – the first to

²²The results on treatment effect dynamics are very similar for the raw samples and are available upon request.

conduct an analysis of the independence property across installations and firms in multiple countries of a large-scale cap-and-trade program targeting greenhouse emissions.

The empirical analysis pursues a difference-in-differences strategy. The causal effect of allowance allocations on emissions is identified by exploiting exogenous variation in installation-level allowance allocations: 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 retain substantial free allocation under a special provision – the 10c rule – for the foreseeable future (control group).

The regression analysis uses a panel of European gas-fired and coal-fired installations 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. I run estimations on raw and matched samples, by matching – within generation technology – on the average of pre-treatment emissions. In addition to the full installation-level sample, I also consider heterogeneity with respect to parent companies' exposure to the EU ETS by separating out installations belonging to multiple-installation firms. Moreover, I conduct the analysis at the firm level.

I fail to reject independence between allocations and emissions in the main analysis and in most sub-samples. Only in the raw firm-level sample that also includes single-installation firms, generally smaller emitters, a statistically significant dependence between allocations and emissions is found. An analysis of treatment effect dynamics shows that this effect is short-lived, as the impact of the change in allocations on emissions reverts back to zero. These results suggest that transaction costs in the EU ETS are fairly low, and mostly fixed rather than variable, as the latter would induce a more permanent dependence between allocations and emissions.

The analysis in this paper has important policy implications, e.g. for adding new sectors to existing cap-and-trade programs 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 administrative and trading costs facing firms are comparable to those for power generators under the EU ETS – policy makers may use free allocation in the political bargaining process without sacrificing the cost-effectiveness of the cap-and-trade program. Of course, whether doing so would be desirable from an equity perspective is an important additional consideration.

One limitation of the paper is that the treatment effect analysis as performed here is

only feasible for the power sector. Of course, cost-effectiveness of the EU ETS may be compromised by emission decisions of installations that are not part of the analysis in this paper, i.e. installations in manufacturing sectors also covered by the program. Further research should investigate this potential channel to losses in cost-effectiveness.

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.
- 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.
- 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.
- Caliendo M. and Kopeinig S. (2008) Some practical guidance for the implementation of propensity score matching. *Journal of Economic Surveys*, 22(1): 31–72.
- 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.
- Dechezleprêtre A., Gennaioli C., Martin R., Muûls M., and Stoerk T. (2019) Searching for carbon leaks in multinational companies.

- 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.
- (2015) Proposal for a Directive of the European Parliament and of the Council Amending Directive 2003/87/EC to Enhance Cost-Effective Emission Reductions and Low-Carbon Investments. *COM(2015) 337 final*.
- 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.

- 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.
- 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.
- 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.
- 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. (2018) Offset credits in the EU ETS: A quantile estimation of firm-level transaction costs. *Environmental and Resource Economics*, 70(1): 77–106.
- 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*.

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.