

Carbon Permits, Plant Emissions and Industry Dynamics

Flora Marchioro

^aInstitute of Economics, Università della Svizzera italiana (USI),

Abstract

Market-based climate policies, such as the EU Emissions Trading System (EU ETS), aim to reduce greenhouse gas emissions while minimizing economic distortions, yet their full impact on firm survival remains debated. While existing research has confirmed the role of the EU ETS in reducing overall emissions, most studies compare regulated and unregulated firms, overlooking variations in policy stringency among ETS-covered firms. Using a difference-in-differences approach, I analyze the emissions of French industrial plants, categorizing them based on permit allocation stringency and pre-existing permit banking. I find that plants subject to stricter permit constraints reduced emissions more than their sectoral peers, but that a portion of these reductions stemmed from plant exits. A survival analysis on industrial and power plants confirms that those facing higher compliance costs due to stricter permit allocation policy were significantly more likely to exit, thereby reshaping industry dynamics. These findings highlight that observed emissions reductions under the EU ETS stem partly from market exits rather than uniform abatement, raising questions about potential impacts of the policy on sector competitiveness.

Keywords: Carbon trading, EU ETS, Allocation rule, Carbon emissions, Plant exit, Industry dynamics

*I thank Massimo Filippini, Sebastien Houde, Raphael Parchet, and Mar Reguant as well as Fabrizio Colella, Patricia Funk, Mario Jametti, Lorenz Küeng, Fabrizio Mazzonna and Giovanni Pica for their helpful discussions. I thank the participants of IdEP Seminar Series, BSE Energy Group Seminar and of IAERE 2025 for their precious feedback. Finally, I thank Jan Abrell, INSEE (Institut National de la Statistique et des Études Économiques) and CASD (Centre d'accès sécurisé distant aux données - CASD project code IMPCLIM) for providing access to the data and continuous technical support. All the results have been reviewed to ensure that no confidential information has been disclosed. All the errors are my own.

1. Introduction

The European Union Emissions Trading System (EU ETS) is the world’s largest carbon market. While extensive research has assessed its impact on emissions and firm outcomes, existing studies have mostly relied on comparisons between regulated and non regulated plants, potentially overlooking important within-policy variation in regulatory stringency. In this sense, evidence on ETS-covered plants exclusively, while accounting for heterogeneity in policy stringency and exposure, is scarce. Moreover, while it is generally confirmed that since its onset the EU ETS led to an overall decline in emissions, the question on whether this reduction stems from firm-level abatement or from compositional effects due to plant exits remains open. This distinction is crucial for evaluating the effectiveness of carbon pricing policies. Indeed, if emissions reductions primarily result from plant closures rather than technological improvements, this raises concerns with respect to industry competitiveness and possible economic distortions.

Addressing these gaps, this paper explores a research question that remains unresolved: Is the reduction in emissions, observed in much of the current literature, due to a general decline across all plants, or is it primarily driven by a compositional effect of the surviving sample?

By developing two measures of policy stringency and policy exposure at the plant level, I analyze the impact that a 2013 change in ETS carbon permit allocation policy had on both plant-level emissions and plant exit rates of French ETS-covered plants. I distinguish between two policy dimensions: first, plants facing higher policy stringency, meaning those with greater free permit constraints relative to their sector median (i.e., dirtier plants); and second, plants with higher policy exposure, meaning those with fewer pre-existing banked permits, making them more vulnerable to compliance costs. I rely on a difference-in-differences (DiD) strategy to examine emissions outcomes in industry plants, defining treatment and control groups based on the policy stringency dimension. I find that plants facing stronger permit policy stringency exhibited greater emission reductions than their competitors within the sector, and that plants began adjusting their emissions levels even before the

policy change was fully implemented. However, the evidence also shows that part of these overall emission reductions can be attributed to plant exits, implying that selection effects played some role in shaping aggregate outcomes. I then develop a survival analysis model for plants subject to different levels of ex-ante policy exposure and policy stringency. Results show that exit risks were significantly higher among plants that faced stricter permit constraints and among plants that were ex-ante more exposed to the policy change, due to their limited reserves of banked permits. Additionally, I find that, also at the sectoral level, policy stringency seems correlated to sectoral exit rates. These findings have important policy implications, as they suggest that the EU ETS not only reduces emissions through plant-level operational adjustments, but also reshapes industry composition.

A large body of literature provides evidence that carbon cap-and-trade mechanisms, and specifically the European Union Emissions Trading System (EU ETS), have successfully reduced overall emissions (e.g. Martin et al. (2014), Martin et al. (2016), Marin et al. (2018), Colmer et al. (2023), Dechezleprêtre et al. (2023)). However, much of this evidence is concentrated on the early phases of the policy (2005–2012), when the system was still in its infancy¹. This temporal limitation is particularly important because, in these early years, the EU ETS was characterized by generous overallocation of free carbon permits to plants, which possibly diluted the effectiveness of the cap-and-trade mechanism. In this respect, my analysis challenges this perspective by arguing that the EU ETS began to have a more substantial impact only in later phases, particularly after 2013, when free allocation rules became stricter and plants were forced to adjust more significantly.

In addition to this temporal limitation, the existing literature has largely relied on a binary comparison of ETS-covered and non-ETS plants, which introduces significant methodological challenges. Many studies match regulated plants with unregulated plants to estimate treatment effects, but this approach may understate the true impact of the policy according to the recent analysis

¹See the literature review provided by Joltreau and Sommerfeld (2016)

by Barrows et al. (2023). Because ETS-covered plants face higher compliance costs, they likely increase prices to compensate for the additional environmental burden². Crucially, non-ETS plants operating in the same output markets also have an interest in raising their prices in response to industry-wide cost pressures. This breaks one of the most fundamental assumptions in difference-in-differences (DiD) analyses, which requires the control group to remain unaffected by the treatment. In the context of existing DiD analyses on the EU ETS, this implies that the policy may have been even more effective than initially thought, potentially driving stronger emission or selection effects (including plant exit) than previously captured³. To address this limitation, I analyze only ETS-regulated plants, classifying them by own developed measures of policy stringency and policy exposure instead of comparing them to non-ETS plants. This approach reduces concerns from Barrows et al. (2023)’s critique, since all plants in my sample face regulatory constraints but with varying compliance costs⁴.

Finally, by focusing on the intensive margin of regulated and unregulated plants, much of the empirical literature mostly studies within-plant operational adjustments while underexploring market-wide compositional effects⁵. If non-ETS plants are indirectly affected and ETS plants pass on costs through pricing, then market structure in certain sectors might itself change, potentially leading to increased plant exits and stronger composition effects than previously estimated. This methodological gap means that previous work may

²Empirical evidence for this is provided by Fabra and Reguant (2014) for the power sector.

³In this regard, and consistent with the first concern, empirical studies that combine pre-post analysis with plant-matching are likely to produce lower-bound estimates. This is because not only they are subject to the critique outlines above, but they also include the effects of the EU ETS during its least stringent phase (Phase I).

⁴A similar approach has been recently applied by De Jonghe et al. (2020).

⁵Many studies have focused on within-firm operational adjustments such as changes in output, investment, R&D, or carbon leakage to non-regulated plants (Martin et al. (2014), Calel (2020), Hintermann et al. (2020), De Jonghe et al. (2020), Dechezleprêtre et al. (2023)). As of my current understanding, Verde et al. (2019) and Guerriero and Pacelli (2023) are out of the few analyses that explicitly study plant entry and exit incentives under the ETS.

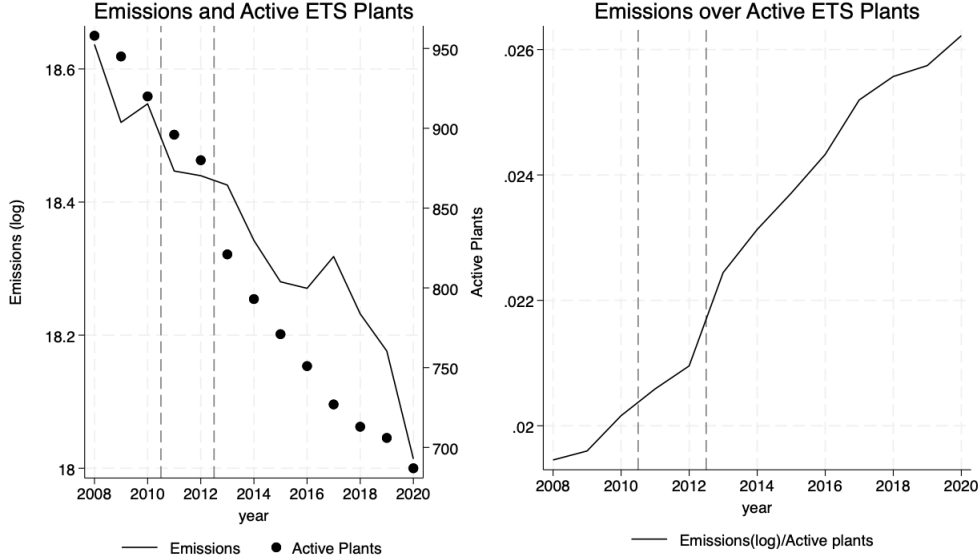
underestimate the EU ETS’s role in shaping market dynamics, firm behavior, and composition of aggregate emissions reductions. By including an analysis on plant survival that relies on policy stringency, I aim to provide a more accurate assessment of how carbon pricing affects emissions and market structure.

Overall, this paper contributes to current gaps in three key ways. First, it introduces a novel approach to assessing within-policy variation in the EU ETS by developing plant-level measures of policy stringency and exposure relative to other plants within the same sector. Second, it provides empirical evidence that a significant share of emissions reductions stems from plant exits rather than uniform abatement across firms. Third, it applies a survival analysis framework to examine how regulatory stringency influences exit probabilities, highlighting sectoral differences and the role of carbon pricing in reshaping industry composition.

To illustrate the potential presence of these compositional effects, Figure 1 provides preliminary evidence on the relationship between emissions trends and number of plants under the EU ETS. The figure shows the evolution of ETS-covered plant-level emissions and the number of active French industry and power plants from 2008 to 2020, based on the sample described in Section 3. The number of active plants and their total verified emissions both decline, with a notable drop around 2013, aligning with stricter EU ETS Phase III regulations. Although both values decline over time, when computing the share of emissions over number of active plants, the ratio between the two appear to consistently increase, suggesting that surviving plants might on average emit more. This trend may indicate that smaller or less efficient plants exited the market, while larger or more competitive plants remained in place.

The paper relates to three main strands of literature. First, the paper connects to the literature on the Coase theorem in cap-and-trade systems Coase (1960), which suggests that the allocation of carbon permits should not affect overall emissions, provided that markets function efficiently. Several studies support this hypothesis, demonstrating that carbon emissions and permit allocations are largely independent due to the flexibility of trading mecha-

Figure 1: Emissions and active French plants covered by the EU ETS



Notes: Active plants include all French plants with positive values of verified emissions in 2008 subject to the EU ETS (i.e. new entrants since 2009 are excluded). The two vertical lines correspond to the years of new allocation rule announcement for EU ETS Phase III (i.e. April 2011) and to its introduction (i.e. April 2013). Data is taken from the European Union Transaction Log (EUTL).

nisms (Reguant and Ellerman (2008), Fowlie and Perloff (2013), Colmer et al. (2023))⁶. This study questions the application of the Coase theorem by examining how permit allocation interacts with firm-level constraints and industry composition effects, potentially leading to asymmetric market outcomes across different markets.

Second, this paper builds on the literature on climate regulation and industry dynamics, which explores how environmental policies influence market composition, exits and acquisitions. Previous research (e.g. Fowlie et al. (2016), Barrows and Ollivier (2018), Verde et al. (2019), De Jonghe et al. (2020), Jo and Karydas (2023)) has established that stringent climate poli-

⁶Studies challenging the Coase hypothesis are mostly focused on evidence in behavioral and managerial economics (Martin et al. (2011), Venmans (2016)) or on corporate finance models (e.g. Bustamante and Zucchi (2022)).

cies can accelerate firm restructuring and output, particularly affecting less competitive or more emissions-intensive firms. The present study extends this literature by focusing on the compositional effects of the EU ETS, examining not only whether firms adapt but also how the regulatory environment influences exit probabilities.

Finally, the paper connects to the literature on carbon trading and corporate finance, which examines plants' responses to carbon permit allocation and trading incentives (Martin et al. (2011), Venmans (2016), Bustamante and Zucchi (2022)). In this regard, my study provides evidence that policy announcement is enough to trigger within-plant adaptation to environmental policy changes, and that plants are differently exposed to permit allocation policy (and, possibly, differently financially-constrained) based on their previous positioning in terms of banking of carbon permits.

The paper is structured as follows. Section 2 provides a context on the EU ETS and on its permit allocation policies across phases II and III. Section 3 outlines an overview of the main data sources and of sample construction. Section 4 presents treatment assignment, as well as the main methodology used to analyze plant emissions and plant exit. Section 5 provides results on plant emissions and plant exits, and comments on them. Finally, Section 6 concludes the analysis.

2. The EU ETS and its Permit Allocation Policy

2.1. The EU Emissions Trading System (EU ETS)

The EU Emissions Trading System (EU ETS) is the largest and most established cap-and-trade program in the world, designed to regulate greenhouse gas emissions (GHGs) from high-emitting sectors such as power generation sector, industry sector and other highly carbon intensive sectors such as waste management⁷. It is characterized by free allocation of non-expiring tradable

⁷The aviation sector was integrated into the EU ETS in 2012, though it is excluded from the scope of this study.

carbon permits given to polluting plants, and a secondary market for transactions of carbon permits (or EU Allowances, EUAs) where polluting plants can buy and sell permits. Established in 2005, the EU ETS covers approximately 12,000 plants across the EU-27, Iceland, Liechtenstein, and Norway.

The EU ETS functions by setting an overall emissions cap that is gradually reduced over time. Plants are required to surrender one EUA for each ton of CO₂ they emit. Those that reduce emissions below their allocation can sell excess permits, while plants exceeding their cap must purchase additional allowances or invest in carbon abatement technologies. This flexibility ensures that emissions reductions occur where they are most cost-effective, aligning with economic principles of market-based regulation (e.g. Coase (1960), Baumol and Oates (1971)). Plants can obtain permits through three primary channels: (i) free allocation from the regulator (i.e. *FA*); (ii) auctioning of permits in the primary auction market, where permits are sold by regulatory authorities; and (iii) trading in the secondary market⁸. Once acquired, plants can use permits in three ways: (1) surrendering them at the end of the compliance year to match their verified emissions; (2) selling excess permits in the secondary market; and (3) banking the excess permits for future use. Banking is permitted across compliance years and trading phases⁹, though borrowing from future periods is prohibited. Specifically in my timespan, permits issued in Phase II and III were non-expiring and could be banked across years. All plants within the established sectors, thresholds and EU Members States are obliged to comply to it, unless they are willing to pay heavy fines on each additional ton of carbon emitted¹⁰.

⁸The secondary market operates through regulated exchanges, such as the European Energy Exchange (EEX), and over-the-counter (OTC) transactions, where buyers and sellers negotiate directly.

⁹The only exception on banking across trading phases was between the first two phases.

¹⁰The annual compliance period in the EU ETS follows a structured cycle to ensure regulated plants surrender enough EUAs to cover their verified emissions. Around February each year, plants receive free allowances and must track their emissions throughout the year. By March 31 of the following year, plants are required to report their verified emissions from the previous year. By April 30, they must surrender a number of EUAs equal to their total emissions. Recall that one EUA corresponds to one ton of CO₂.

The EU ETS has evolved through distinct trading phases, each introducing refinements to the allocation mechanism and the scope of regulation. Phase I (2005-2007) served as a pilot phase with generous permit allocation and volatile carbon prices. Since this phase was known since its onset to be a pilot phase (and thus its policy stringency was limited), since Phase II permits were canceled after the end of the phase, and since my analysis uses at least partly information on the amount of banked permits at the plant level, including this phase does not add anything to my analysis. Phase II (2008-2012), instead, introduced stricter caps and limited auctioning but retained the reliance on National Allocation Plans (NAPs), just like the pilot phase. Phase III (2013-2020) marked a fundamental shift, eliminating NAPs, introducing benchmark-based allocation for industry, and implementing full auctioning for power plants.

2.2. Permit Allocation Policies


2.2.1. Phase II (2008-2012)

During Phase II (2008-2012), permit allocation was governed by National Allocation Plans (NAPs), under which each Member State set its own allocation rules. In France, for example, allocations were based on 2005 plant-level emissions data and projected sector growth for the 2008-2012 period¹¹. The allocation policy during this phase raised several concerns. First, as documented in Rogge et al. (2006), allowing each EU Member State to determine carbon allocation rules for its own plants led to excessive permit distributions and raised concerns about potential political influence from industrial groups. Second, evidence of windfall profits in highly concentrated sectors emerged, with plants able to more than pass through the implicit cost of permits to consumers while continuing to receive free allocations (as documented by Fabra and Reguant (2014) for the power sector). Third, the 2008 financial crisis exacerbated permit overallocation. As industrial production declined, emissions in some sectors fell below sector-level projected growth rates included in NAPs,

¹¹ Allocation rules based on historical emission levels are referred to by Fowlie et al. (2016) as pure grandfathering.

and thus resulted in an increased surplus of unused permits. Evidence of over-allocation of permits to polluting plants is presented in Figure 2, based on the sample of French industry and power plants described in Section 3. Overallocation is measured at the plant level as the ratio of free allocated permits over verified emissions for each plant in the sample, and then collapsed into yearly averages. A value above 1 means that plants on average received more free allowances than needed to cover their verified emissions, while a value below 1 means plants received fewer allowances than their emissions, requiring them to purchase additional permits, use their banked permits, or reduce emissions to comply. The figure illustrates how plants in Phase II (2008-2012) received more free allowances than necessary, with the average ratio of allocated permits to verified emissions well above 1 and possibly leading to a surplus of banked permits. The reasons for the steep decline after 2013, different for the sample of industry plants only vs the sample including power plants, is presented in the following section.

Figure 2: Overallocation: plant free allocated permits over plant emissions



Notes: Calculation based on the sample described in Section 3. Calculation does not include permits auctioned from the primary market, nor carbon offsets. The decrease starting in 2017 is likely due to the announcement of the Market Stability Reserve (MSR) mechanism.

2.2.2. Phase III (2013-2020)

Partly due to concerns over market distortions under excessive overallocation, windfall profits, and pressure from domestic industrial groups on national regulators, the European Commission amended its permit allocation rules for Phase III (2013-2020). Permit allocation rules for the ETS Phase III were broadly based on two main European Commission Directives: EC (2009) of 23 April 2009 (hereafter, 2009 Directive), and EC (2011) of 27 April 2011 (hereafter, 2011 Directive).

The 2009 Directive introduced distinct allocation rules for the power and non-power sectors, with significantly different levels of clarity regarding their future obligations¹². Article 10 of the directive explicitly mandates that from 2013 onwards, power generators must participate in full auctioning of emissions allowances, thus explicitly eliminating free allocation for these plants. The rationale behind this decision is that power companies can pass on the cost of emissions allowances to consumers through electricity prices Fabra and Reguant (2014). Article 10a further clarifies that electricity generators would not receive any free allocation except in specific cases, such as district heating or high-efficiency cogeneration. This unambiguous policy direction meant that power plants had somehow full certainty regarding their future compliance obligations, allowing them to start adapting as early as 2009.

In contrast, the situation for industry plants (i.e., all other non-power sector installations) was much less clear. While Article 10a mentions that transitional free allocation would be provided, in contrast to NAPs used so far, it did not define specific allocation rules. Instead, the article states that the European Commission would develop harmonized allocation principles through future implementing measures. Overall, then, the directive set broad objectives, such as the intention to use ex-ante benchmarks based on the average performance of the 10% most efficient installations in each sector during 2007-2008 (so-called, benchmarking), but did not specify how these benchmarks

¹²The European Commission defines an electricity generator as an installation that, on or after 1 January 2005, has produced electricity for sale to third parties.

would be applied across different industries nor did it present any rule to build these benchmarks¹³. This shift aimed to incentivize emissions efficiency by rewarding best practices.

Based on this distinction in both policy stringency (i.e. full auctioning vs. benchmarking) and the clarity of policy guidelines, power plants likely began reacting to the 2013 changes immediately after the directive announced in 2009, while industrial plants might have postponed their response until further clarification. Such clarification on the newly-adopted benchmarking rule was only provided with the 2011 Directive EC (2011)¹⁴. Overall, given the complexity of the rule and the postponement of its finalization up until 2011 only, industrial plants remained uncertain until 2011 about the final allocation rules. Unlike power plants, which had clarity by 2009, industry plants faced a shorter adjustment period to prepare for their sharp 2013 changes to full auctioning. For the reasons outlined in this and the next section, I carefully consider my treatment to start either in 2011 (i.e. policy announcement under non-power plants) or under 2013 (i.e. policy implementation for all plants). Additionally, due to their inherent differences in policy announcement timing (2009 vs 2011) and in intensity of treatment (full auctioning vs benchmarking), I refer to power plants vs non-power plants (hereafter, industry plants) and analyse them separately¹⁵.

3. Data Sources and Sample Overview

3.1. Data Sources

The present study combines two main sources of data. First, plant-level emission data from the EUTL (or *European Union Transaction Log*). The

¹³Additionally, Article 10a(5) introduces a cross-sectoral correction factor (CF), suggesting potential adjustments to free allocation, but without clear details on how it would impact individual sectors.

¹⁴More information on benchmarking as outlined in the 2011 Directive can be found in Appendix A.

¹⁵As outlined in the next section, since the non-power, non-industry sector is a relatively small percentage of my sample (7%) and is subject to the same treatment as the industry sector, I include it as part of the industry sector

EUTL is the official registry of the EU ETS and it provides a list of all regulated installations, past and present. Second, firm-level data are gathered by the French statistical institute, INSEE (or *Institut National de la Statistique et des Études Économiques*). Specifically, INSEE offers two main databases used in this analysis: FICUS-FARE (or the unification of *Fichier Complet Unifié de SUSE* until 2007 and *Fichier Approché des Résultats d'ESANE* since 2008); and EACEI (or *Enquête annuelle sur les consommations d'énergie dans l'industrie*). The fiscal census FARE-FICUS offers annual income statements and balance sheets of the universe of French firms in manufacturing, mining, utilities and service sector.

Data on ETS covered plants are obtained from the EUTL, as processed by Abrell (2021). Plants in this database are recorded in terms of compliance information (i.e. number of yearly free allocated permits, surrendered permits, verified emissions, and daily transactions of permits in the primary and secondary market), registered in terms of city, postalcode, geographical coordinates and NACE-4 digits, as well as connected to their respective account owner. In turn, each plant-connected account owner records a company registration number, which in the French case coincide with French firm identifiers (i.e. SIREN or *Système d'Identification du Répertoire des Entreprises*). Account holders can then be identified as the firms owning ETS plants.

Accordingly, these firms can hence be mapped with the FARE-FICUS and EACEI data provided by INSEE. FARE-FICUS provides general information about the firm (SIREN identifier, industry classification, head office address, total number of workers employed, age, etc.), the income statement (containing variables such as total turnover, total labour costs, and value added) as well as balance sheet information (e.g. various measures of capital, debt, and assets). While starting from a list of plant-level identifiers, i.e. SIRET codes (or *Système d'Identification du Répertoire des Etablissements*), one could trace back to the firm level databases¹⁶, the opposite is not true if one possesses the

¹⁶Indeed, the plant SIRET identifier is a 14 digit number whose first 9 digits correspond to the SIREN firm identifier.

SIREN code only (as in my case). However, to identify the number of plants owned by a firm in a specific year, I then briefly take advantage of the annual employment database at the plant level (or DADS, *Déclarations Annuelles des Données Sociales*). From this database I identify how many SIRET plant codes are connected to the same SIREN firm codes each year, and I consider this as the number of plants owned by a single firm in a specific year. From this same database, I also keep the geographical location of plants owned by a firm. I am then able to match the ETS plants with the EACEI survey (covering industry plants only), based on the SIREN code and geographical location of plants in the survey¹⁷. From the EACEI survey, I observe quantities and values of energy consumed by fuel type (i.e. electricity, steam, fossil fuels, and biofuels). As in Jo and Karydas (2023), I aggregate the consumption of different sources of energy to a clean and a dirty bundle for each plant, with the clean bundle including electricity, steam and renewables and the dirty bundle consisting of all other fuels (natural gas, petroleum products, etc.).

3.2. Sample Overview

The main sample is constructed in the following way. The EUTL initial sample of French plants outside of the aviation sector is composed of 1,542 plants. Out of these, 78 exited the ETS before 2013 and 415 entered the sample after 2008¹⁸. The sample therefore drops to 1,046 plants. Additionally, I exclude individual plants that are registered in the EUTL but that never registered positive values of carbon emissions, resulting in a sample of 880 plants. Finally, as presented in the next section, due to the structure of the treatment, the sample on which treatment assignment can be defined relies on the shock between 2012 and 2013 in permit allocation to be negative, i.e. a *drop* in allocated permits. I hence exclude from the analysis 117 plants

¹⁷To avoid mismatching of firms owning more than one plant within the same postal code, I match ETS plants to energy surveyed plants only when the SIREN-postal code pair in the DADS database is one only. In other words, I avoid matching through SIREN firms that own multiple SIRET plants within the same postal code.

¹⁸A plant is considered closed in the analysis if it records zero surrendered emissions in the EUTL for two consecutive years.

that were active before 2013 and experienced an *increase* in allocated carbon permits, instead. The final EUTL sample is then composed of 762 plants (295 power plants, 467 industry plants) owned by 531 unique firms, which are followed yearly from 2008 to 2020.

Out of this EUTL sample, using the SIREN code of the firm I match 490 of these 531 firms to FARE-FICUS data on the universe of French firms (or around 717 plants, 278 power plants and 439 industry plants). Out of these industry plants, I match 261 of them with the EACEI plant-level energy survey data.

A summary statistics of the available variables for the full sample is presented in Table 1. Emission levels verified by the regulator are very similar to surrendered emissions by plants (that is, surrendered permits), confirming that plant compliance to the policy is high. Most importantly, these values are overall higher than the value of allocated free permits, FA , confirming that in the overall sample, pre-2013 and post-2013, the free allocation policy was binding for plants. However, permit banking is on average higher than surrendered or emissions, supporting the idea that overallocation of permits is a concern in the overall sample (pre-post 2013 and for all plants). Power plant firms account for overall 38% of the full sample, industry plants account for (54%) and other sectors (e.g. waste management) account for only 7%. Turning to firm-level data, ETS plants are connected to firms that own on average more than 35 other plants, and that are considered quite sizeable in terms of employment and fixed assets. Employment at the plant level is measured as firm-level employment over the number of plants owned by the firm. For the subset of industry plants only, energy survey data is matched on a combination of SIREN firm code and postal codes of the surveyed plant, as outlined in the previous section. Energy intensity at the surveyed plant level is measured as plant-level energy consumption, over firm-level output sold.

Table 1: Summary of data sources and variables

	All plants, sectors, and years	
	Mean	SD
EUTL: Plant-level data		
Allocated permits	99.51	504.06
Verified emissions	108.99	495.38
Surrendered emissions	109.08	497.78
Banking of permits	113.41	464.39
Power sector	0.38	0.49
Observations	9906	
FARE-FICUS: Firm-level data		
Employment (firm)	4844.52	15388.95
Fixed assets	4953.74	20284.73
Output sold	2392.27	8430.12
Nr. plants	36.98	194.84
Employment (plant)	367.51	852.76
Observations	9321	
EACEI: Plant-level energy survey		
Clean energy consumption	79.55	169.25
Dirty energy consumption	149.19	306.98
Energy consumption	228.73	383.56
Energy intensity	2.91	39.22
Observations	3393	

Notes: The sample used here includes all plants (treated and controls), all sectors (industry and power) and all years (pre and post treatment). Allocated permits, verified emissions, surrendered permits and actual net banking are expressed in thousands of EU carbon permits. Fixed assets and production sold are expressed in thousands of Euros. Energy consumption variables are expressed in thousands, where "clean" is composed of the sum of electricity and steam, while "dirty" is composed of coal, oil and natural gas. Energy intensity is measured as energy consumption over output sold.

4. Methodology and Treatment Assignment

4.1. Treatment Assignment

The main treatment variable identifies plants that experienced a greater-than-median reduction in freely allocated permits within their sector. The main treatment variable, $high_drop_i$, captures differences in policy stringency across plants and is defined as follows ¹⁹:

$$high_drop_i = \begin{cases} 1 & \text{if } drop_2013_i \geq drop_2013_{s,median}, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

$$\text{where } drop_2013_i = -\frac{FA_{i,2013} - FA_{i,2012}}{FA_{i,2012}}.$$

Here, $FA_{i,t}$ represents the quantity of free allocated permits to plant i in year t . A plant is classified as $high_drop_i$ if its reduction in free permits be-

¹⁹This treatment assignment resembles the one in De Jonghe et al. (2020).

tween 2012 and 2013 was at least as severe as the median reduction observed within its sector. This variable thus reflects the relative policy stringency imposed on a plant, compared to other plants in the same sector. Since plant-level free allocated permits are still compared to their previous level when building the $drop_2013_i$ variable, and plant-level drop is only later compared to other drops within the sector, this approach effectively takes into account both the plant-level component and the sector-level benchmarks outlined in Section 2. Indeed, plants that experience high $drop_2013_i$ compared to their $FA_{i,2012}$ values (i.e. or equally compared to their $FA_{i,2008}$ values, see Appendix A) can be considered to be *far* away to the EU-sector benchmark of emission efficiency²⁰. Subsequently, plants that are then also assigned to the $high_drop_i$ treatment can be considered *further* away than EU-sector benchmark of emission efficiency than their French competitors within the sector. In other words, $high_drop_i$ treated plants could be interpreted as dirtier plants compared to both their EU and French counterparts. Overall, the treated group presenting higher-than-median drops in free allocated permits in 2013 is composed of 236 industry plants, while the control group is composed of 231 plants. A summary statistics and pre-treatment balance test is presented in Table 2. Although the two groups appear to differ with respect to certain plant-level and firm-level variables, the addition of plant-level fixed effects should limit concerns in this regard.

To explore heterogeneity in plants' ability to respond to the reform, plants are further categorized based on their banking behavior in the years prior to the policy shift. A secondary variable, $high_bank_i$, is used to analyze heterogeneous effects by distinguishing plants based on their pre-existing stock of banked permits prior to the policy change:

²⁰A distribution of the $drop_2013_i$ variable by sector is reported in Figure Appendix B.1.

Table 2: Balance test of control vs treated group, Industry plants, Pre-2011

	Low Drop		High Drop		Diff (b)
	Mean	SD	Mean	SD	
Plant-level data					
Allocated permits	229.37	947.80	134.77	406.06	94.61*
Verified emissions	195.25	806.18	109.34	376.23	85.91*
Surrendered emissions	195.28	806.18	109.36	376.23	85.93*
Permits annual banking	246.81	1251.07	154.43	506.09	92.38
Employment (plant)	218.32	380.55	480.07	1162.71	-201.75**
Firm-level data					
Employment (firm)	1409.47	2643.19	3438.64	16670.24	-2029.18**
Fixed assets	714.90	1275.76	1067.16	3606.58	-352.26*
Output sold	493.34	691.45	740.39	1937.90	-247.05**
Nr. plants	10.53	19.02	7.96	13.42	2.57*
Energy data (plant-level)					
Clean energy consumption	79.05	135.29	95.70	219.35	-16.65
Dirty energy consumption	108.44	106.26	191.67	444.94	-83.23*
Energy consumption	187.49	206.41	287.37	537.69	-99.88**
Energy intensity	1.72	2.16	9.15	123.21	-7.43
Observations	693		708		1401

Notes: 2011 is kept as reference year due to possible policy anticipation at the 2011 announcement. Allocated permits, verified emissions, surrendered permits and actual net banking are expressed in thousands of EU carbon permits. Fixed assets and production sold are expressed in thousands of Euros. Energy consumption variables are expressed in thousands, where "clean" is composed of the sum of electricity and steam, while "dirty" is composed of coal, oil and natural gas. Energy intensity is measured as energy consumption over output sold. A t-test measuring the difference between the control and treated groups is presented in the last column.

$$high_bank_i = \begin{cases} 1 & \text{if } bank_{i,08-10} \geq bank_{s,08-10median}, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

$$\text{where } bank_{i,08-10} = \frac{1}{3} \sum_{t=2008}^{2010} bank_{i,t}.$$

The variable $bank_{i,08-10}$ represents the average quantity of permits banked by plant i between 2008 and 2010²¹. Plants classified as $high_bank$ had a 2008-2010 stock of banked permits that was at least as large as the median within their sector over this period. Compared to their sector competitors,

²¹Recall that permits across Phases II and III of the EU ETS are non-expiring and can be banked both across compliance years and across compliance phases. However, borrowing from subsequent periods is not allowed. The $bank_{i,t}$ variable is then a cumulative flow variable defined at plant i and year t which takes into account: (1) the non-negative amount of $FA_{i,t}$ to a plant; (2) the non-negative amount of banked permits at $t - 1$; (3) the positive or negative net trading of permits performed by the plant at year t (i.e. permit auctioning from the primary market, and permit purchases or sales in the secondary market) (4) the non-negative amount of permits the plant has to surrender at the end of the compliance year in line with how much carbon it emitted.

these plants are considered better positioned to mitigate the compliance costs associated with the reduction in freely allocated permits in 2013, as they could draw on previously accumulated allowances. This variable thus reflects the relative policy exposure of a plant, compared to other plants in the same sector.

As shown in Figure 3, when looking at plant surrendered emissions classified by permit drop and banking group variables in the industry sample, both groups of low- and high-banking plants (in black and blue, respectively) behave similarly to policy announcement in 2011 and to policy introduction in 2013, although they differ in terms of level of emissions²². Although high-banking plants emit more than low-banking ones, it appears there is no strong difference between the two groups with respect to reaction to policy announcement. For this reason, in the main difference-in-difference analysis on industry plant surrendered emissions in Section 4.2.1 I do not sample the *high_drop_i* treatment differently according to the banking assignment, but instead bundle the two banking groups together. In other words, I compare the *low_drop* plants to the *high_drop* plants, irrespective of the sampling based on *high_bank_i* variable. A similar emission graph sampled only by *high_bank_i* variable is found separately in Figure Appendix B.2.

Turning to plant exit behavior based on *high_drop_i* and *high_bank_i* assignment in Figure 4, most exits since 2013 appear to involve plants experiencing a high drop in allocated free permits compared to their sector median. Additionally, the groups corresponding to *high_bank* (in blue) register fewer exits than the low-banking group. Most importantly, as the sample is by definition restricted to plants active until 2013, exits in the pre-period are zero for all groups, which prevents a different-in-different analysis on the outcome of plant exit. Instead, as outlined in the next section, I will rely on survival analysis using both *high_drop_i*, *high_bank_i* and their interaction as treatment variables. A similar exit graph for power plants is found separately in Appendix B.4.

²²See Appendix B.3 for the emission graph on power sector plants.

Figure 3: Plant emissions, industry sector, by drop and banking groups





Figure 4: Plant exit, industry sector, by drop and banking groups



Before turning to the identification strategy and results, it is important to stress a key difference between the samples analysing the two main outcomes. On the one hand, when analysing the plant emissions the sample excludes power plants due to anticipated permit policy announcement (i.e. April 2009), as outlined in Appendix A. Indeed, in case of policy anticipation (which I later show to be present for industry plants, at least), a difference-in-difference analysis including power plants would necessarily require to set a baseline year before 2009, thus preventing any possible pre-period analysis. Additionally, this sector is the only one that was defaulted to full permit auctioning since 2013, on top of being clearly informed about it since 2009 already. In this sense, including power plants to the analysis of the other sectors might make the interpretation of average treatment estimates more difficult. Hence, this sector was excluded from the analysis on plant emissions. On the other hand, the identification through survival analysis on plant exit does not rely on a pre-treatment period and, by sample definition, would consider the outcome of plant exit just after 2013, regardless of policy announcement. In this sense, main results for plant exit also include the power sector, although I will show overall results are robust to its exclusion.

4.2. Identification Strategy

4.2.1. Plant Emissions

To estimate the effects of carbon policy stringency on plant-level outcomes, a difference-in-differences (DiD) model is implemented using the following specification:

$$\log Y_{i,t} = \sum_{k=2008, k \neq b}^{2020} \beta_k \text{high_drop}_i \times \mathbb{1}[t = k] + \alpha_i + \lambda_t + \theta X_{i,t} + \epsilon_{i,t} \quad (3)$$

where $b \in \{2010, 2012\}$

where $Y_{i,t}$ represents the main outcome of interest for plant i at time t (i.e. plant surrendered emissions). The coefficients β_k measure the estimated effect of being in the *high_drop_i* treatment group relative to the baseline

year. To account for possible policy anticipation before 2013, I indeed estimate two different models based on 2010 or 2012 as baseline year. The term α_i captures plant fixed effects, absorbing time-invariant plant-specific characteristics. The coefficients λ_t include year fixed effects to account for common time trends. The vector $X_{i,t}$ represents additional firm-level control variables (not included in the results presented here), while $\epsilon_{i,t}$ is the plant-clustered error term. The interaction terms between $high_drop_i$ and year dummies capture the relative effect of policy stringency on treated plants (i.e. experiencing an above-median drop in free allocated permits) versus untreated plants (i.e. experiencing a below-median drop in free allocated permits), with 2010 as the baseline reference year before policy announcement.

4.2.2. Plant Exit

To analyze the probability of plant exit over time, I also estimate a Cox proportional hazard model:

$$h_i(t) = h_0(t) \times \exp \left(\beta_1 high_drop_i + \beta_2 high_bank_i + \right. \quad (4)$$

$$\left. \beta_3 (high_drop_i \times high_bank_i) + \gamma_s + \epsilon_i \right)$$

where $h_i(t)$ represents the hazard function which measures the instantaneous probability of exit for plant i at time t , given that it has survived up to that point. The baseline hazard function, $h_0(t)$, remains unspecified in the Cox model and serves as a reference for estimating relative risks. The main explanatory variables include $high_drop_i$ and $high_bank_i$, as presented above. The interaction term between the two accounts for differential effects when both conditions are met. Additionally, I control for unobserved heterogeneity across industries by including sector fixed effects, γ_s . I estimate this model using maximum likelihood estimation (MLE), and interpret the Cox coefficients β_k in terms of hazard ratios²³.

²³Recall that the relationship between Cox estimates and Hazard ratios is Hazard Ratio = $\exp(\beta_k)$. A hazard ratio greater than 1 indicates an increased likelihood of plant exit, while a hazard ratio less than 1 suggests a lower likelihood of plant exit.

5. Results

5.1. *Results on Plant Emissions*

Regression estimates using 2010 as baseline year are presented in Table 3. Treatment estimates for the full sample (column 1) display a statistically significant decline in emissions for high-drop plants relative to low-drop plants starting in 2013, with larger reductions persisting in later years. These results confirm that the observed emissions reductions in the graphical analysis can be attributed to the policy rather than to underlying trends in emissions. Moreover, the estimated effects after 2013 are particularly strong among plants with low banking levels, suggesting that plants with fewer banked permits were more constrained in their ability to adjust once policy is in place. This evidence aligns with the expectation that plants with lower compliance flexibility would experience stronger emissions reductions in response to the allocation policy change.

I then estimate the same model using 2012 as the baseline year (Table B.4). Unlike the 2010-based estimates, the 2012-based results show no significant difference in emissions reductions between high- and low-drop plants after policy introduction. This divergence between the two specifications highlights a key empirical challenge: plants may have anticipated the policy and adjusted their emissions behavior before its implementation in 2013. By using 2010 as the baseline, I indeed capture pre-policy introduction responses that might be missed when using 2012 as the reference year. Overall then, while Figure 3 would lead to believe that plant emissions started decreasing since policy introduction only, the DiD estimates suggest that plants might have begun adjusting their behavior at policy announcement already.

To understand how robust these results are to sample composition and plant exits, I separately run the main regression of column 1 of Table 3 for the sample of plants that remained active throughout the entire sample (i.e. 2008-2020). Results are presented in Table 4. Compared to full sample estimates, when restricting the analysis to active plants only the estimated effects become weaker and in several cases lose statistical significance. The comparison suggests that, rather than all plants adjusting their emissions downward

Table 3: DiD estimates on industry plant emissions

	(1) Full Sample	(2) Low Bank	(3) High Bank
2008.year1.high_drop	0.016 (0.034)	-0.008 (0.035)	0.039 (0.059)
2009.year1.high_drop	0.050 (0.035)	0.028 (0.047)	0.072 (0.053)
2010.year1.high_drop	0.000 (.)	0.000 (.)	0.000 (.)
2011.year1.high_drop	-0.046 (0.036)	-0.032 (0.038)	-0.059 (0.061)
2012.year1.high_drop	-0.127*** (0.046)	-0.093* (0.055)	-0.159** (0.072)
2013.year1.high_drop	-0.142*** (0.054)	-0.119* (0.067)	-0.163* (0.084)
2014.year1.high_drop	-0.154** (0.069)	-0.156 (0.100)	-0.153 (0.097)
2015.year1.high_drop	-0.143** (0.063)	-0.170* (0.088)	-0.123 (0.090)
2016.year1.high_drop	-0.223*** (0.070)	-0.234** (0.093)	-0.214** (0.102)
2017.year1.high_drop	-0.211*** (0.069)	-0.224** (0.095)	-0.200** (0.099)
2018.year1.high_drop	-0.182** (0.071)	-0.126 (0.088)	-0.224** (0.108)
2019.year1.high_drop	-0.216*** (0.072)	-0.183** (0.092)	-0.241** (0.109)
2020.year1.high_drop	-0.199** (0.082)	-0.202* (0.118)	-0.194* (0.116)
Observations	5534.000	2591.000	2943.000
R-squared	0.942	0.915	0.924

Standard errors in parentheses

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

Notes: 2010 taken as baseline year. Plant and year fixed effects included in all models. The addition of sector fixed effects do not change estimates nor standard deviations. Standard errors clustered at the plant level.

Table 4: DiD estimates on industry plant emissions, by plant status

	(1)	(2)	(3)
	Full Sample	Active only	Exiting only
2008.year1.high_drop	0.016 (0.034)	-0.011 (0.041)	0.136* (0.072)
2009.year1.high_drop	0.050 (0.035)	0.066 (0.041)	-0.020 (0.120)
2010.year1.high_drop	0.000 (.)	0.000 (.)	0.000 (.)
2011.year1.high_drop	-0.046 (0.036)	0.021 (0.039)	-0.170* (0.093)
2012.year1.high_drop	-0.127*** (0.046)	-0.045 (0.049)	-0.319** (0.129)
2013.year1.high_drop	-0.142*** (0.054)	-0.053 (0.052)	-0.483** (0.199)
2014.year1.high_drop	-0.154** (0.069)	-0.069 (0.064)	-0.444 (0.306)
2015.year1.high_drop	-0.143** (0.063)	-0.087 (0.066)	-0.281 (0.219)
2016.year1.high_drop	-0.223*** (0.070)	-0.129* (0.069)	-0.784** (0.324)
2017.year1.high_drop	-0.211*** (0.069)	-0.147** (0.073)	-0.488** (0.243)
2018.year1.high_drop	-0.182** (0.071)	-0.128* (0.076)	-0.205 (0.268)
2019.year1.high_drop	-0.216*** (0.072)	-0.159** (0.074)	0.206 (0.856)
2020.year1.high_drop	-0.199** (0.082)	-0.146* (0.085)	
Observations	5534.000	4835.000	699.000
R-squared	0.942	0.946	0.892

Standard errors in parentheses

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

Notes: 2010 taken as baseline year. Plant and year fixed effects included in all models. The addition of sector fixed effects do not change estimates nor standard deviations. Standard errors clustered at the plant level.

in response to the policy, a significant portion of the observed decline is attributable to plant exits. This compositional effect is confirmed when looking at the estimates for exiting plants only in column 3: despite the much smaller sample size, the treatment effects are large, strongly negative, and statistically significant, indicating that exiting plants might be a key driver in full sample estimates.

Figure 5 reports the estimates for the first two columns in Table 4 ²⁴. The results illustrate that prior to 2013, the estimated coefficients remain close to zero, confirming that both samples followed similar pre-trends. However, post-2013, we observe a divergence: the estimated effects for the full sample (which includes exiting plants) become increasingly negative, indicating a significant reduction in emissions for high-drop plants relative to low-drop plants. The effects for active plants follow a similar trend but are generally smaller in magnitude, suggesting that part of the emissions reduction observed in the full sample is driven by plant exits rather than reductions within continuing plants.


5.2. Results on Plant Exit

Results on plant exit are presented in Table 5. The hazard ratio of 1.097 for *high_drop* in column 1 implies that plant experiencing a high drop in free permits face a 109.7% higher exit risk compared to those with a lower-than-median drop. Conversely, the hazard ratio of 0.501 for *high_bank* in column 2 suggests that plants with high banking levels before 2011 are around 50% less likely to exit, reinforcing the idea that higher banking acts as a buffer against exit risk. Results are robust to adding both variables simultaneously, as well as adding sector fixed effects (columns 3 and 4).

Finally, looking at the interaction between *high_drop* and *low_bank* (i.e. the inverse of *high_bank*) in column 5, one can compare the most policy affected and most exposed plants to the rest of the sample. The results confirm the expected pattern: *high_drop* alone significantly increases the hazard ratio

²⁴Figure Appendix B.5 report instead the estimates for exiting plants only.

Figure 5: Plant emissions DiD Estimates, industry sector, Full sample vs Active Plants only



for exit, and *low_bank* alone also increases exit risk. However, the interaction term is negative and not statistically significant, meaning that the combined effect of *high_drop* and *low_bank* does not appear to significantly amplify exit risk beyond their individual contributions.

Overall, these results suggests that, while plants with higher permit banking levels before 2011 are more resilient to the permit allocation shock, this attenuating effect is not strong enough to fully offset the impact of a sharp reduction in free allocated permits. Plants with high free permit shock remain significantly more likely to exit, even if they had previously accumulated permits.

Table 6 shows instead the results when limiting the sample to industry plants only. Compared to the full-sector analysis, the effects of *high_drop* and *high_bank* remain consistent in direction but appear more pronounced. More specifically, while in the sample including power plants the estimates for *high_drop* were higher than in analysis of industry plants only, the opposite

Table 5: Survival Analysis: Cox Proportional Hazard Models

	(1)	(2)	(3)	(4)	(5)
High Drop	1.097*** (0.168)		1.191*** (0.169)	1.302*** (0.172)	
High Bank		-0.501*** (0.152)	-0.660*** (0.153)	-0.781*** (0.157)	
High Drop = 1					1.428*** (0.317)
Low Bank = 1					0.921*** (0.334)
High Drop \times Low Bank					-0.183 (0.383)
Sector FE	No	No	No	Yes	Yes
Observations	762.000	762.000	762.000	762.000	762.000
LR Chi-sq	48.208	11.022	67.095	127.072	127.306
Log-likelihood	-1138.500	-1157.093	-1129.056	-1099.068	-1098.951

Standard errors in parentheses

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

Table 6: Survival Analysis: Cox Proportional Hazard Models (Industry plants only)

	(1)	(2)	(3)	(4)	(5)
High Drop	1.016*** (0.228)		1.030*** (0.228)	1.157*** (0.231)	
High Bank		-1.208*** (0.236)	-1.221*** (0.236)	-1.310*** (0.238)	
High Drop = 1					1.969*** (0.619)
Low Bank = 1					2.106*** (0.615)
High Drop \times Low Bank					-1.009 (0.673)
Sector FE	No	No	No	Yes	Yes
Observations	467.000	467.000	467.000	467.000	467.000
LR Chi-sq	22.232	30.436	53.307	103.853	106.488
Log-likelihood	-564.369	-560.268	-548.832	-523.559	-522.242

Standard errors in parentheses

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

is true for *high_bank*. This is in line with the hypothesis that when including power plants in the analysis, for which the permit policy change in 2013 was stronger, *high_drop* is a stronger predictor of plant exit and not enough attenuated by the amount of banked permits. In other words, policy stringency dominates the protective effect of permit banking. Instead, when focusing the analysis to industry plants only, the two opposite effects *high_drop* and *high_bank* appear to attenuate each other more: pre-existing permit stocks might average play a more crucial role in preventing industry plants from exiting.

Table 7: Median Permit Drop, Cox Estimates and Hazard Ratios, by Sector

Sector	Median Permit Drop	Cox Estimate	Hazard Ratio
Chemicals & Pharma (5)	-0.2711	-1.022***	0.360***
Metals (6)	-0.2916	-1.152*	0.316*
Food & Beverages (2)	-0.3655	-0.779***	0.459***
Electronics (7)	-0.3701	-34.842	7.39E-16
Textiles & Wood & Paper (3)	-0.3740	-0.446	0.640
Coke & Petrol & Plastics (4)	-0.3959	-0.192	0.825
Machinery & Transportation (8)	-0.4245	0.592	1.807
Water & Waste (11)	-0.4635	-34.887	7.06E-16
Trade & Transport (13)	-0.4950	-0.024	0.976
Services (14)	-0.4950	1.213***	3.363***
Agriculture & Mining (1)	-0.6871	1.226*	3.407*

Standard errors in parentheses

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

Notes: Sector classification from 1 to 14 is based on the NACE 2-digits assigned to the plant in the EUTL database. Sector 10 (Power) is taken as baseline. Sector 9 (Other Manufacturing) and sector 12 (Construction) are dropped due to missing observations. Sectors 7 (Electronics) and 11 (Water and Waste) that experienced zero exits in the observed period correctly present hazard ratios close to zero. Sectors are ordered from the highest to lowest sector median drop (column 1). Hazard ratios in column 3 are computed based on the Cox estimates (column 2) of the hazard model with sector fixed effects in Table 5, column 4. For the full model see Appendix B.5

Finally, the results presented in Table 7 report the sector fixed effect hazard rate coefficients of the last column in Table 5, along with sector median drops. This table highlights strong differences in exit probabilities across sectors, using the power sector (i.e. sector 10) as baseline. The decision on using power sector is supported both by the fact that only this sector experienced full auctioning since 2013 (see Section 2), and because this sector appears to

consistently have higher exit rates than all others combined, as evident from Appendix B.6. Results are ordered decreasingly in terms of sector median drops. Comparing sector median drops (interpreted as how stringent reductions of free permits were to plants in that sector) and sector hazard ratios, there appears to be a correlation between policy stringency on the sector and likelihood of plant exit at the sector level. For example, the services sector and the agriculture and mining sector, which have some of the largest median permit drops, also display the highest hazard ratios, indicating a significantly increased probability of exit relative to the baseline. Conversely, sectors with lower median permit drops (i.e. with plants less strongly affected by the 2013 policy change), such as chemicals and pharma, metals, and food and beverages, exhibit lower hazard ratios.

6. Conclusions

This paper examines the impact of a 2013 change in the free permit allocation rule for plants within the European Union Emissions Trading Scheme (EU ETS) on both plant-level carbon emissions and plant exit dynamics. While previous studies have documented emissions reductions under the EU ETS, much of the existing literature focuses on earlier phases of the EU cap-and-trade, when the policy lacked sufficient stringency. Additionally, no prior research has analyzed the impact of the ETS beyond the simple distinction between regulated and unregulated plants, and specifically exploiting variation in permit allocation within the sample of ETS-regulated plants. Finally, evidence on potential compositional effects within this group remains extremely scarce. This study contributes to filling these gaps by analyzing whether the observed emissions reductions since 2013 resulted from an overall decline in emissions across all plants, or were instead driven by a shift in emissions due to plant exits.

Using a difference-in-differences approach, I analyze the effect of the policy change in free allocated plant permits on plant emissions. I distinguish industry plants along two lines: plants that faced stronger policy stringency compared to their sector median (i.e. the supposedly dirtier plants), and plants

that were more exposed to the policy due to a limited stock of pre-existing banked permits, whose presence could have attenuated the policy impact. The results indicate that emissions reductions started well before policy implementation, and confirm that stricter permit allocation rules play a role in driving plant-level emission reductions. However, these reductions were not evenly spread across surviving plants but were instead concentrated among those that exited the market. Overall, this suggests that a portion of the observed decline in emissions was driven by plant closures rather than within-plant abatement efforts.

Additionally, I apply a survival analysis to assess the likelihood of plant exit in response to the policy stringency. The survival analysis further confirms the role of plant exit in shaping aggregate emissions trends. Dirtier plants within their sector — or more precisely plants that experienced a more severe reduction in free permit allocations — in both the power and non-power sectors, record a significantly higher probability of exit, whereas those with higher initial permit banks appear more resilient. Additionally, the overall pattern suggests that the sector-level policy stringency, as proxied by the median permit drop, is an important determinant of plant survival for most sectors excluding those that experience no plant exit. This appears to support the hypothesis that sectors subject to stricter environmental policy (in terms of lower amount of free carbon permits) may have been more likely to experience plant exit, confirming the compositional effect in emissions reductions.

These findings have important implications for the design of emissions trading schemes. While this study confirms that the EU ETS effectively reduced overall emissions, it is the first to provide empirical evidence that compositional effects contributed to these reductions, with a portion of the decline being driven by plant closures rather than exclusively by technological improvements or process efficiencies. This highlights the need for policymakers to carefully consider how allocation mechanisms and compliance flexibility measures influence firm behavior, particularly in industries like the power sector already characterized by high market concentration, facing high capital costs, or subject to limited short-term abatement options. Still, the present

research has not yet been able to account for plant- or firm-level output and market share consideration needed to provide a more comprehensive overview of the issue, and would thus greatly benefit from a more structural approach to predict plant exit. Future research in this direction could further explore the extent to which plant exits contributed to output changes, sectoral restructuring and market concentration in response to carbon pricing policies.

References

- Abrell, J. (2021). Database for the european union transaction log (eutl.info). Technical report.
- Barrows, G., Calel, R., Jégard, M., and Ollivier, H. (2023). Estimating the effects of regulation when treated and control firms compete: A new method with application to the eu ets. CESifo Working Paper Series 10438, CESifo.
- Barrows, G. and Ollivier, H. (2018). Cleaner firms or cleaner products? how product mix shapes emission intensity from manufacturing. *Journal of Environmental Economics and Management*, 88:134–158.
- Baumol, W. J. and Oates, W. E. (1971). The use of standards and prices for protection of the environment. *Swedish Journal of Economics*, 73:42–54.
- Bustamante, M. C. and Zucchi, F. (2022). Dynamic carbon emission management. *SSRN Electronic Journal*.
- Calel, R. (2020). Adopt or innovate: Understanding technological responses to cap-and-trade. *American Economic Journal: Economic Policy*, 12(3):170–201.
- Coase, R. H. (1960). The problem of social cost. *The Journal of Law Economics*, 3:1–44.
- Colmer, J., Martin, R., Muûls, M., and Wagner, U. J. (2023). Does Pricing Carbon Mitigate Climate Change? Firm-Level Evidence From the European Union Emissions Trading Scheme. CRC TR 224 Discussion Paper Series crctr224_2023_32v2, *University of Bonn and University of Mannheim, Germany*.

- De Jonghe, O., Mulier, K., and Schepens, G. (2020). Going green by putting a price on pollution: Firm-level evidence from the EU. *SSRN Electron. J.*
- Dechezleprêtre, A., Nachtigall, D., and Venmans, F. (2023). The joint impact of the european union emissions trading system on carbon emissions and economic performance. *Journal of Environmental Economics and Management*, 118:102758.
- EC (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 (text with eea relevance). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32009L0029>.
- EC (2011). 2011/278/eu: Commission decision of 27 april 2011 determining transitional union-wide rules for harmonised free allocation of emission allowances pursuant to article 10a of directive 2003/87/ec of the european parliament and of the council (notified under document c(2011) 2772). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32011D0278>.
- EC (2013). 2013/448/eu: Commission decision of 5 september 2013 concerning national implementation measures for the transitional free allocation of greenhouse gas emission allowances in accordance with article 11(3) of directive 2003/87/ec of the european parliament and of the council (notified under document c(2013) 5666). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex>
- Fabra, N. and Reguant, M. (2014). Pass-through of emissions costs in electricity markets. *American Economic Review*, 104(9):2872–99.
- Fowlie, M. and Perloff, J. (2013). Distributing pollution rights in cap-and-trade programs: Are outcomes independent of allocation? *The Review of Economics and Statistics*, 95(5):1640–1652.
- Fowlie, M., Reguant, M., and Ryan, S. P. (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*, 124(1):249–302.

- Guerriero, C. and Pacelli, A. (2023). Emissions Abatement: the Role of EU ETS and Free Allowances. The Italian Case. CSEF Working Papers 698, Centre for Studies in Economics and Finance (CSEF), University of Naples, Italy.
- Hintermann, B., Žarković, M., Maria, C. D., and Wagner, U. J. (2020). The Effect of Climate Policy on Productivity and Cost Pass-Through in the German Manufacturing Sector. CRC TR 224 Discussion Paper Series crctr224_2020_49, *University of Bonn and University of Mannheim, Germany*.
- Jo, A. and Karydas, C. (2023). Firm heterogeneity, industry dynamics and climate policy. *SSRN Electronic Journal*.
- Joltreau, E. and Sommerfeld, K. (2016). Why does emissions trading under the EU ETS not affect firms' competitiveness? Empirical findings from the literature. Technical report.
- Marin, G., Marino, M., and Pellegrin, C. (2018). The Impact of the European Emission Trading Scheme on Multiple Measures of Economic Performance. *Environmental & Resource Economics*, 71(2):551–582.
- Martin, R., Muûls, M., and Wagner, U. (2011). Climate change, investment and carbon markets and prices—evidence from manager interviews. *Climate Strategies, Carbon Pricing for Low-Carbon Investment Project*, 2011:1–51.
- Martin, R., Muûls, M., de Preux, L. B., and Wagner, U. J. (2014). Industry compensation under relocation risk: A firm-level analysis of the eu emissions trading scheme. *American Economic Review*, 104(8):2482–2508.
- Martin, R., Muûls, M., and Wagner, U. J. (2016). The impact of the european union emissions trading scheme on regulated firms: What is the evidence after ten years? *Review of Environmental Economics and Policy*, 10(1):129–148.
- Reguant, M. and Ellerman, A. D. (2008). Grandfathering and the endowment effect - an assessment in the context of the spanish national allocation plan. Working papers, Massachusetts Institute of Technology, Center for Energy and Environmental Policy Research.

- Rogge, K. S., Schleich, J., and Betz, R. (2006). An early assessment of national allocation plans for phase 2 of eu emission trading. urn:nbn:de:0011-n-494736.
- Sartor, O., Pallière, C., and Lecourt, S. (2014). Benchmark-based allocations in eu ets phase 3: An early assessment. *Climate Policy*, 14.
- Venmans, F. M. J. (2016). The effect of allocation above emissions and price uncertainty on abatement investments under the eu ets. *Journal of Cleaner Production*, 126:595–606.
- Verde, S. F., Graf, C., and Jong, T. (2019). Installation entries and exits in the eu ets: patterns and the delay effect of closure provisions. *Energy Economics*, 78:508–524.

Appendix A. Permit Allocation Rule in ETS Phase III (2013-2020) for industry plants

According to the 2011 Directive EC (2011), and as modeled in Sartor et al. (2014), Phase III allocation for non-electricity generation plants was built based on the following formula:

$$FA_{i,j,t} = B_j \cdot P_{i,j} \cdot RF_{j,t} \cdot CF_t \quad (\text{A.1})$$

where $FA_{i,j,t}$ represents the free carbon permits allocation over products j produced in plant i in year t . The total allocation $FA_{i,t}$ is obtained by summing all $FA_{i,j,t}$ for all products j produced in plant i in year t . The term B_j (i.e. sectoral-level component of the policy) refers to 54 product benchmarks, which were built on the basis of average emissions of the 10% most emission-efficient plants in 2007-2008 at the EU level. The benchmark values were determined using the arithmetic average of the greenhouse gas performance of the most efficient installations during this period, as specified in Annex I of EC (2011). The parameter $P_{i,j}$ (i.e. plant-level component of the policy) corresponds to the highest median historical production of product j by plant i in either the period 2005-2008 or 2009-2010, depending on which is higher, as specified in Annex III of EC (2011). The term $RF_{j,t}$ represents the reduction factor in allocations applied to products that are not at risk of carbon leakage. This factor decreases linearly from 0.8 to 0.3 between 2013 and 2020 in order to prevent overallocation. The list of sectors s exempt from this reduction factor due to carbon leakage risk is detailed in Annex VI of EC (2011). Finally, the correction factor CF_t is applied uniformly across sectors s , decreasing from 0.94 to 0.82, to ensure that the total free allocation does not exceed the maximum emissions cap, as required by Annex II EC (2013).

An example of free allocation of carbon permits by plants in different sectors is displayed in A.1. Plant A belongs to the power sector while Plant B belongs to the industry sector (based on EUTL-defined NACE 4-digits classifications). As mentioned in the main text, grandfathering (G) in Phase II was based on French National Allocation Plan (NAPs) allocated permits based on

Table A.1: Allocation example with a power and an industry plant

Year	Free Permits A	Allocation A	Sector A	Free Permits B	Allocation B	Sector B
2008	107,750	G	Steam	103,050	G	Textiles
2009	107,750	G	Steam	103,050	G	Textiles
2010	107,750	G	Steam	103,050	G	Textiles
2011	107,750	G	Steam	103,050	G	Textiles
2012	107,750	G	Steam	103,050	G	Textiles
2013	0	A	Steam	59,192	B	Textiles
2014	0	A	Steam	58,164	B	Textiles
2015	0	A	Steam	57,123	B	Textiles
2016	0	A	Steam	56,072	B	Textiles
2017	0	A	Steam	55,010	B	Textiles
2018	0	A	Steam	53,937	B	Textiles
2019	0	A	Steam	52,851	B	Textiles
2020	0	A	Steam	51,761	B	Textiles

Notes: Plant A belongs to the power sector; Plant B belongs to the industry sector. Grandfathering (G) in Phase II was based on French National Allocation Plan (NAPs) allocated permits based on a mix of 2005 emissions and projected industry growth. Auctioning (A) corresponds to no zero permits allocated. Benchmarking (B) in Phase III was based on EU-wide best practices. The rule was approved in 2011 for all years of Phase III.

a mix of 2005 emissions and projected industry growth. Auctioning (A) corresponds to no zero permits allocated. Benchmarking (B) in Phase III drops in from 2012 to 2013 due mostly to the B_j and $P_{i,j}$ components, and then decreases yearly due to the $RF_{j,t}$ and CF_t components.

Appendix B. Additional Graphs and Tables

Table B.1: Summary statistics of Pre-2011 vs Post-2011, Industry plants

	Pre-2011		Post-2011		Diff
	Mean	SD	Mean	SD	
Plant-level data					
Allocated permits	181.56	727.69	125.71	585.85	55.85**
Verified emissions	151.84	628.16	121.89	558.17	29.95
Surrendered emissions	151.86	628.15	122.03	558.23	29.83
Permits annual banking	200.13	951.38	132.50	311.38	67.63**
Employment (plant)	379.76	872.08	379.49	894.78	0.27
Firm-level data					
Employment (firm)	2419.10	11945.43	2175.14	9086.97	243.96
Fixed assets	890.32	2705.53	1221.21	3581.83	-330.88***
Output sold	616.37	1457.32	986.35	3161.05	-369.98***
Nr. plants	9.25	16.51	11.02	19.09	-1.77**
Energy data (plant-level)					
Clean energy consumption	87.42	182.48	78.89	173.88	8.53
Dirty energy consumption	150.28	326.63	156.25	317.79	-5.97
Energy consumption	237.71	410.78	235.14	395.94	2.57
Energy intensity	5.43	87.13	2.28	15.91	3.15
Observations	1401		4670		6071

Notes: Allocated permits, verified emissions, surrendered permits and actual net banking are expressed in thousands of EU carbon permits. Fixed assets and production sold are expressed in thousands of Euros. Energy consumption variables are expressed in thousands, where "clean" is composed of the sum of electricity and steam, while "dirty" is composed of coal, oil and natural gas. Energy intensity is measured as energy consumption over output sold.

Table B.2: Summary statistics of treated vs controls, Industry plants

	Low Drop		High Drop		Diff (b)
	Mean	SD	Mean	SD	
Plant-level data					
Allocated permits	195.63	830.56	82.78	289.48	112.84***
Verified emissions	182.85	765.35	75.90	275.16	106.94***
Surrendered emissions	182.87	765.34	76.10	275.39	106.78***
Permits annual banking	169.96	687.94	129.92	378.00	40.04**
Employment (plant)	350.08	816.56	410.21	960.38	-60.13
Firm-level data					
Employment (firm)	1514.15	3150.52	2967.35	13561.79	-1453.20***
Fixed assets	1005.55	2041.56	1288.43	4376.05	-282.88**
Output sold	862.41	2995.24	941.30	2724.91	-78.89
Nr. plants	11.89	21.62	9.37	14.86	2.52***
Energy data (plant-level)					
Clean energy consumption	74.33	122.70	87.01	217.25	-12.68
Dirty energy consumption	127.98	156.83	183.35	425.76	-55.37***
Energy consumption	202.31	236.10	270.36	514.17	-68.05***
Energy intensity	1.62	2.26	4.25	58.91	-2.63
Observations	3003		3068		6071

Notes: Allocated permits, verified emissions, surrendered permits and actual net banking are expressed in thousands of EU carbon permits. Fixed assets and production sold are expressed in thousands of Euros. Energy consumption variables are expressed in thousands, where "clean" is composed of the sum of electricity and steam, while "dirty" is composed of coal, oil and natural gas. Energy intensity is measured as energy consumption over output sold.


Table B.3: Summary statistics of treated vs controls, Industry plants, Post-2011

	Low Drop		High Drop		Diff (b)
	Mean	SD	Mean	SD	
Plant-level data					
Allocated permits	185.50	791.96	67.19	241.82	118.31***
Verified emissions	179.12	752.81	65.87	235.75	113.25***
Surrendered emissions	179.15	752.80	66.12	236.10	113.03***
Permits annual banking	143.48	307.95	120.48	314.73	23.01*
Employment (plant)	366.76	887.05	392.90	903.08	-26.14
Firm-level data					
Employment (firm)	1544.54	3283.11	2827.22	12490.69	-1282.68***
Fixed assets	1091.27	2210.43	1355.18	4581.55	-263.91*
Output sold	971.26	3380.19	1001.90	2918.66	-30.65
Nr. plants	12.31	22.34	9.78	15.23	2.53***
Energy data (plant-level)					
Clean energy consumption	73.23	119.62	84.86	216.78	-11.62
Dirty energy consumption	132.52	166.10	181.28	421.06	-48.77***
Energy consumption	205.75	242.43	266.14	508.34	-60.39***
Energy intensity	1.59	2.29	3.01	22.74	-1.42*
Observations	2310		2360		4670

Notes: Allocated permits, verified emissions, surrendered permits and actual net banking are expressed in thousands of EU carbon permits. Fixed assets and production sold are expressed in thousands of Euros. Energy consumption variables are expressed in thousands, where "clean" is composed of the sum of electricity and steam, while "dirty" is composed of coal, oil and natural gas. Energy intensity is measured as energy consumption over output sold.

Figure Appendix B.1: Distribution of $drop_2013_i$ variable, by Sectors

Distribution of drop_2013 by Sector



Notes: Sector classification from 1 to 14 is based on the NACE 2-digits assigned to the plant in the EUTL database. Sector 9 (Other Manufacturing) and sector 12 (Construction) are dropped due to missing observations.

Figure Appendix B.2: Plant emissions, industry sector, by drop groups




Figure Appendix B.3: Plant emissions, power sector, by drop and banking groups




Figure Appendix B.4: Plant exit, power sector, by drop and banking groups




Figure Appendix B.5: Plant emissions DiD Estimates, industry sector, Exiting plants only




Table B.4: Industry plant emissions, 2012 baseline

	(1) Full Sample	(2) Low Bank	(3) High Bank
2008.year1.high_drop	0.143*** (0.048)	0.085 (0.060)	0.198*** (0.073)
2009.year1.high_drop	0.177*** (0.044)	0.121** (0.059)	0.231*** (0.067)
2010.year1.high_drop	0.127*** (0.046)	0.093* (0.055)	0.159** (0.072)
2011.year1.high_drop	0.081*** (0.031)	0.061 (0.046)	0.100** (0.043)
2012.year1.high_drop	0.000 (.)	0.000 (.)	0.000 (.)
2013.year1.high_drop	-0.015 (0.040)	-0.026 (0.063)	-0.004 (0.050)
2014.year1.high_drop	-0.027 (0.061)	-0.063 (0.102)	0.005 (0.073)
2015.year1.high_drop	-0.017 (0.056)	-0.077 (0.086)	0.036 (0.074)
2016.year1.high_drop	-0.096 (0.064)	-0.141 (0.089)	-0.055 (0.091)
2017.year1.high_drop	-0.084 (0.062)	-0.131 (0.091)	-0.041 (0.085)
2018.year1.high_drop	-0.055 (0.066)	-0.033 (0.086)	-0.065 (0.095)
2019.year1.high_drop	-0.089 (0.068)	-0.090 (0.098)	-0.082 (0.094)
2020.year1.high_drop	-0.072 (0.079)	-0.109 (0.125)	-0.035 (0.103)
Observations	5534.000	2591.000	2943.000
R-squared	0.942	0.915	0.924

Standard errors in parentheses

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

Figure Appendix B.6: Exit rates, by sector

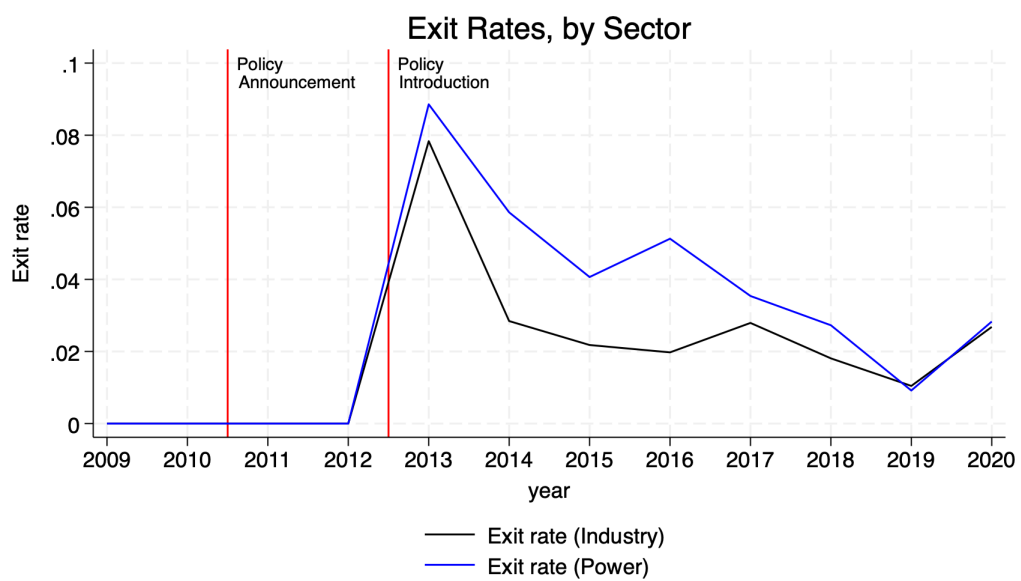


Table B.5: Survival Analysis: Cox Proportional Hazard Models (Industry Hazard Rates)

	(1)
high_drop	1.302*** (0.172)
high_bank	-0.781*** (0.157)
1.sector	1.234* (0.717)
2.sector	-0.708** (0.281)
3.sector	-0.453* (0.274)
4.sector	-0.199 (0.395)
5.sector	-1.019*** (0.238)
6.sector	-1.160** (0.589)
7.sector	-34.813 (4.63e+07)
8.sector	0.589 (0.394)
10.sector	0.000 (.)
11.sector	-34.860 (2.97e+07)
13.sector	-0.048 (0.395)
14.sector	1.213*** (0.282)
Observations	762.000
LR Chi-sq	127.072
Log-likelihood	-1099.068

Standard errors in parentheses

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

Notes: Sector 10 (Power) taken as baseline. Sector 9 (Other Manufacturing) and Sector 12 (Construction) are dropped due to missing observations.