

Does Pricing Carbon Mitigate Climate Change? Firm-Level Evidence from the European Union Emissions Trading Scheme*

Jonathan Colmer
University of Virginia

Ralf Martin
Imperial College London

Mirabelle Muûls
Imperial College London

Ulrich J. Wagner
University of Mannheim

November 2020

Abstract

In theory, market-based regulatory instruments correct market failures at least cost. However, evidence on their efficacy remains scarce. We evaluate the European Union Emissions Trading Scheme (EU ETS) – the world’s first and largest market-based climate policy. Using administrative data on almost 4,000 French manufacturing firms, we estimate that the EU ETS induced regulated firms to reduce carbon dioxide emissions by 8-12% compared to unregulated firms after the Pilot phase, a necessary condition for climate change mitigation. These reductions account for 26% of the concurrent decline in aggregate industrial emission in France. We do not estimate any negative effects on the scale of production; instead we find that firms reduced the emissions intensity of value added by making targeted investments. We find no evidence that firms outsourced production to unregulated firms or markets. Collectively, these findings suggest that the EU ETS induced global emissions reductions, a necessary and sufficient condition for mitigating climate change.

JEL codes: Q54 Q58, H23, L50, F18

*Correspondence to m.muuls@imperial.ac.uk. We thank Michael Best, Meredith Fowlie, Andreas Gerster, Michael Greenstone, Stephen Holland, Mar Reguant, Johannes Spinnewijn, and John Van Reenen for very helpful thoughts and discussions. We are also grateful to seminar participants at Cambridge, ETH Zurich, Geneva, LSE, Mannheim, NYU Shanghai, Oxford, Toulouse, and Waseda, as well as to conference participants at the NBER Summer Institute, TSE, IZA, EUI Florence School of Regulation, and at several other conferences for helpful comments and suggestions. All the results have been reviewed to ensure that no confidential information has been disclosed (Centre d’accès sécurisé distant aux données, CASD - project E598). Funding by the Economic and Social Research Council (ESRC) through grant no. ES/J006742/1, Centre for Climate Change Economics and Policy, the Grantham Foundation, European Union’s Horizon 2020 research and innovation programme under grant agreement no. 308481, the German Research Foundation (DFG) through CRC TR 224 (Project B7), is gratefully acknowledged. This work is also supported by a public grant overseen by the French National Research Agency (ANR) as part of the “*Investissements d’Avenir*” program (reference: ANR-10-EQPX-17 – *Centre d’accès sécurisé distant aux données – CASD*) All errors and omissions are our own.

1 Introduction

The unchecked accumulation of greenhouse gas (GHG) emissions is one of the starkest examples of market failure worldwide. GHG emissions are a by-product of valuable economic activities. However, the costs they impose through climate change are not fully accounted for in economic decision-making. In theory, market-based regulations hold the promise of mitigating climate change at least cost to society (Pigou, 1920; Baumol & Oates, 1971; Baumol, 1972; Montgomery, 1972; Tietenberg, 1973; Nordhaus, 1977; Hahn, 1989; Nordhaus, 2001; Burke et al., 2016; Gillingham & Stock, 2018).¹ These regulations discourage the production of emissions-intensive goods by putting a price on emissions. The price encourages both emissions abatement, in particular by emitters with low abatement costs, and investments in technology that lowers abatement costs.

Market-based regulations allow polluting firms more flexibility in choosing their own path to compliance than command-and-control regulation, yet different compliance strategies have very different implications for the economy and the global environment. Flexibility in how to comply may lead to leakage effects that undermine climate change mitigation. If regulated firms cut emissions by outsourcing carbon-intensive elements of the value chain to unregulated firms or markets, then carbon emissions will simply ‘leak’ to unregulated jurisdictions. Carbon leakage threatens the efficacy of any unilateral climate change mitigation policy by limiting, or even reversing, its impact on global emissions.

This paper provides evidence on the environmental and economic consequences of market-based regulations to mitigate climate change by evaluating the European Union Emissions Trading Scheme (EU ETS) – the world’s first and largest market-based climate policy. Introduced in 2005, the EU ETS establishes a price for the right to emit carbon dioxide (CO₂) emissions. This is achieved by imposing a cap on the aggregate emissions from more than 12,000 power and manufacturing plants in 31 countries. The cap covers 45% of EU emissions and 5% of global emissions. Tradeable permits are then issued for each tonne of CO₂ under the cap. The permit price is formed in a European wide market where firms with a permit surplus sell to firms that require permits in order to comply with the regulation.

Whether such a cap-and-trade scheme reduces emissions is a question of regulatory strin-

¹While there is plenty of disagreement among economists in discussions of policy and government intervention, a preference for market-based regulatory instruments is a point in which economists largely agree. On January 17th 2019, over 3,500 economists, from a diverse set of political, ideological, and academic backgrounds, rallied around the efficacy of market-based mechanisms for internalizing the social costs of climate change in a statement published in the Wall Street Journal – the largest public statement by economists in history. The second largest public statement by economists was the “Economists’ Statement on Climate Change” signed by 2,500 economists in 1997 at the time of the Kyoto Protocol, calling for market-based mechanisms to mitigate climate change.

gency. That is, emissions within the regulated market must be lower than if the cap did not exist. In lieu of this unobservable condition, economists view a high and stable permit price as a credible signal of regulatory stringency. During Trading Phase I (from 2005 until 2007), permit prices first climbed to \$37 before collapsing to less than \$1 in early 2007. However, permit prices rebounded to around \$21.35 (\$2017) during Trading Phase II (between 2008 and 2012). Whether these prices were sufficient to deliver meaningful reductions in emissions is an empirical question. We answer this question based on comprehensive administrative data from the French manufacturing sector.

Using a difference-in-differences research design, we estimate that the EU ETS induced regulated firms to reduce CO₂ emissions relative to unregulated firms by 8-12%, during Trading Phase II, but not in Trading Phase I when prices were more volatile. We estimate no significant effects prior to the announcement of the EU ETS or during the announcement period. On aggregate, our results imply that CO₂ emissions fell by 4.7 million tonnes each year, accounting for approximately 26% of France’s aggregate industrial emissions reductions during this period.

We also provide evidence to suggest that the EU ETS induced global emissions reductions, which is the relevant outcome from the perspective of climate change mitigation. First, we estimate no detectable negative effects on the economic performance of regulated firms. If we found such effects, this could mean that the policy shifted production and emissions to unregulated firms. Counter to this leakage mechanism, we estimate significant reductions in CO₂ emitted per Euro of value added, but no effect on value added or employment. Second, we find no evidence that firms increased imported inputs or the carbon content of inputs through trade. Nor do we estimate increased substitution towards purchased electricity. These findings are inconsistent with carbon leakage being a first-order driver of the estimated emissions reductions in this context. Rather, we present evidence that investing in cleaner production processes was the prevailing abatement mechanism among regulated firms.

Overall, our findings suggest that the EU ETS induced regulated firms to reduce CO₂ emissions by reducing the emissions intensity of production with no detectable negative effects on economic performance or outsourcing to unregulated markets. The maximum permit price during the time of the estimated emissions reductions suggests that marginal abatement costs could not have exceeded \$53 per tonne of CO₂ (\$2017). This price reflects the point where firms would have been indifferent between buying permits and reducing emissions and so true marginal abatement costs were likely much lower. Nevertheless, this cost compares favorably to the marginal abatement costs of many non-market based regulatory instruments (Gillingham & Stock, 2018). To the degree that these insights generalize to other markets and settings, our study highlights that market-based regulations can, in

practice, be an effective and economically reasonable tool for mitigating climate change.

Our paper contributes to several literatures. First, we contribute to a literature exploring the effects of environmental regulation on firm behavior (Becker & Henderson, 2000; Greenstone, 2002; Fowlie et al., 2012; Greenstone et al., 2012; Ryan, 2012; Walker, 2013; Martin et al., 2014a,b; Fowlie et al., 2016; He et al., 2020). This literature typically focuses on the effects of policy on either economic or environmental outcomes. We evaluate treatment effects on both types of firm-level outcomes. We also provide detailed evidence on the mechanisms through which firms reduce emissions. As explained above, this is essential if we are to understand whether the policy was effective at achieving its ultimate objective, which is to reduce global emissions.

Second, we contribute to a growing empirical literature seeking to understand the effects of the EU ETS itself (Martin et al., 2016). Early studies in this area have been at the country or sector-level, which complicates causal inference due to confounding factors (Ellerman & Feilhauer, 2008; Ellerman et al., 2010; Egenhofer et al., 2011; Andersen & Di Maria, 2011). Most relevant to our study is a strand of the literature that employs difference-in-differences designs akin to Fowlie et al. (2012) in order to evaluate the impacts of the EU ETS on manufacturing firms.² A robust finding across studies is the absence of detrimental effects on economic performance, broadly defined (Petrick & Wagner, 2014; Jaraite & Di Maria, 2016; Klemetsen et al., 2020; Marin et al., 2018; Dechezleprêtre et al., 2018). The available evidence on industrial CO₂ emissions is not conclusive, however, and results vary across countries and trading phases. Specifically, emissions reductions were estimated for Norway (Klemetsen et al., 2020) but not for Lithuania where CO₂ intensity fell (Jaraite & Di Maria, 2016). The EU ETS was found to have no impact on CO₂ intensity in the United Kingdom (Calel, 2020), though it may have reduced CO₂ emissions in that country, according to a study of selected emitters in four EU countries (Dechezleprêtre et al., 2018). These studies are valuable because they establish under which conditions the EU ETS induced local reductions in emissions. The principal limitation in previous research is a lack of compelling evidence on the mechanisms through which emissions reductions were delivered. Yet understanding the mechanisms is crucial if we are to rule out the possibility that local emissions reductions did not translate into global reductions, which is a necessary and sufficient condition for mitigating climate change. Our study fills this gap. Using linked administrative data from multiple sources, not only do we estimate the effects of the EU ETS on the emissions and economic performance of firms, but we also identify how firms respond to comply with the

²Beyond manufacturing, researchers have estimated the impact of the EU ETS on power plants (Fabra & Reguant, 2014; Zaklan, 2020), on patenting (Calel & Dechezleprêtre, 2016), and on foreign direct investment (Koch & Basse Mama, 2019; Borghesi et al., 2020).

regulation. In so doing, we provide the first evidence in support of the proposition that the EU ETS, the most significant climate policy instrument to date, has delivered on its stated policy objective.

Finally, we provide early empirical evidence that market-based mechanisms are a cost-effective way of reducing emissions. In recent years there has been renewed interest in understanding which government interventions are most effective at improving social welfare (Hendren & Sprung-Keyser, 2020; Hendren & Finkelstein, *ming*). Such questions are especially important in the context of mitigating global climate change, due to the severity of the problem and due to the limited resources available to tackle it. The emissions reductions induced by the EU ETS likely cost substantially less per tonne of CO₂ than alternative non-market-based regulatory instruments (Gillingham & Stock, 2018).

In the next section, we describe the design of the EU ETS and our empirical approach. Section 3 presents the data used for analysis. Section 4 presents the main results and explores the underlying mechanisms. Section 5 present back-of-the-envelope calculations that consider the contribution of the EU ETS to aggregate emissions reductions and compares the cost-effectiveness of the EU ETS to other existing and proposed climate change mitigation policies. Section 6 concludes.

2 Evaluating the European Union Emissions Trading Scheme

Identifying the causal effects of a real-world policy intervention is never a trivial exercise. In the context of the EU ETS, two major challenges arise. First, accurate data on carbon emissions prior to the implementation of the ETS is scarce, as most countries did not explicitly collect this information before it was required for monitoring purposes.³ However, pre-implementation data is required to establish that any measured change in the performance of regulated firms can plausibly be ascribed to the policy itself, and not to other factors. With access to rich administrative data on the fuel use of French manufacturing plants, we are able to construct a consistent, bottom-up measure of direct emissions for all firms, including unregulated ones, both before and after the implementation of the EU ETS. Each dataset as well as the linkages are explained in detail in Section 3 below.

Second, to evaluate the effects of any policy, it is important to have a credible counterfactual. This is particularly challenging in the absence of experimental conditions in which

³Previous work on this policy has been largely unable to compare emissions before and after its introduction Ellerman & Buchner (2008); Ellerman et al. (2010); Egenhofer et al. (2011); Andersen & Di Maria (2011).

subjects can be randomly assigned to treatment and control groups. Correlation does not imply causation. There are many reasons why emissions could have fallen since the implementation of the EU ETS. Emissions in Europe have been declining for some time, as a result of structural economic change and due to energy efficiency improvements. Furthermore, the Great Recession caused economic activity to drop significantly, which in turn led to a further drop in greenhouse gas (GHG) emissions in the EU and around the world. These trends make the evaluation of emissions trading schemes at the aggregate level (i.e., country or sector) a futile exercise, because it is not possible to disentangle the effects of policy changes from other changes over time.

Therefore, it is only through the combination of temporal and cross-sectional variation in treatment assignment among otherwise similar firms that one can hope to identify the causal effect of the EU ETS on emissions and economic outcomes. The remainder of this section explains why the design of the EU ETS gives rise to both types of variations and how the specific institutional details allow us to identify and estimate the causal impact of the policy using variants of the difference-in-differences estimator.

2.1 Treatment Assignment in the EU ETS

The EU ETS is a European wide cap-and-trade program for CO₂ emissions.⁴ Polluters regulated under the policy are required to surrender, at the end of each year, one European Union Allowance (EUA) for each tonne of gas they have emitted over the year. They may buy additional EUAs or sell excess EUAs on an international market at a uniform price. Within limits, EUAs can be banked or borrowed to balance needs across years. The total amount of EUAs in the system is limited and linearly declines over time. Scarce EUAs command a positive price in the permit market. The treatment effect we seek to identify is the average effect of having to pay for CO₂ emissions on various outcome variables of treated polluters.⁵

Our identification strategy exploits both temporal and cross-sectional variation in treatment assignment. The EU ETS was launched in 2005, when France and most other European countries did not have CO₂ prices in place. We thus consider the year 2005 as the beginning of the treatment period. In addition, we allow for the possibility that polluters

⁴Ellerman et al. (2016) provide a concise yet comprehensive review of the history and structure of the EU ETS.

⁵Allocation of EUAs to polluters is via free allocation or permit auctions, with a trend from the former to the latter. During the study period of this paper, however, free permit allocation to manufacturing firms was the rule. Our main analysis abstracts from permit allocation for two reasons. First, by a Coasian argument, permit allocation should not affect firm behavior at the margin. Second, we lack a credible strategy to test for a causal effect. We do provide evidence, however, that differences in the initial allocation of permits give rise to heterogeneous treatment effects (see Section 4.4).

responded to the announcement of the policy before the actual launch.⁶ The EU ETS was officially announced with the publication of the Emissions Trading Directive in 2003 (Directive 2003/87/EC). However, the publication of the directive marked the culmination of a multi-year consultation process between the EU Commission and stakeholders about key design features of the policy. The process was initiated with the publication of a green paper by the EU Commission in 2000 (European Commission, 2000). Comments on the green paper submitted by businesses, NGOs and governments were published in May 2001 (European Commission, 2001). At that point, actors likely had some clarity regarding the shape that the ETS would be taking. We thus consider the year 2001 as the beginning of the announcement period.

Cross-sectional variation in treatment assignment arises because not all CO₂ emitters in Europe are regulated under the EU ETS. Participation criteria were first spelled out in the Emissions Trading Directive and then transposed into national laws.⁷ These criteria are targeted at industrial facilities at the sub-plant level, referred to in the directive as installations. Different criteria are defined for combustion activities on one hand and other carbon intensive processes on the other hand.

In France, participation is mandatory for combustion installations with a rated thermal input of 20 megawatts (MW) or more. This concerns fossil-fuel fired power plants but also industrial plants across a wide range of industries which generate heat, steam or power on site. Further industrial installations are included because they specialize in carbon intensive processes and exceed specific capacity thresholds. Process-based definitions target, *inter alia*, pulp and paper mills, coke ovens, petroleum refineries, non-metallic mineral products (including the manufacture of glass, ceramics, and cement), and the manufacture of basic metals.⁸ Indirect emissions, i.e. from sources that are not owned and not directly controlled by the firm, are not taken into account, nor are electricity imports.

As we explain in further detail in Section 3.7 below, we match French ETS installations listed in the official trading registry to the manufacturing establishments operating them. Any establishment identified in this way is considered as treated and referred to as an ETS plant. Likewise, a firm is considered as treated and referred to as an ETS firm if it owns at least one ETS plant.

The installation-centered, capacity-based participation rules used in the Emissions Trad-

⁶Since CO₂ intensities are often embodied in long-lived capital goods, such anticipated adjustments make economic sense if they prevent a polluter from being locked into high CO₂ intensities – and hence, high compliance costs – for decades to come.

⁷To harmonize criteria across countries, as well as to include additional sectors, the directive was later amended (Directive 2009/29/EC)

⁸Beginning in 2012, emissions from other industries, such as aviation, have been included in the ETS as well.

ing Directive induce variation in treatment status even among plants and firms of similar size (Petrick & Wagner, 2014; Calel & Dechezleprêtre, 2016). To see this, consider as an example the case of two plants that operate combustion installations, as depicted in Figure 1. Both plants have a total combustion capacity of 30 MW, but the distribution of that capacity across installations gives rise to different treatment assignments. Plant 2 is treated because it operates a *single* installation with a rated thermal input of 30 MW, which is above the participation threshold. Plant 1 is not regulated because it achieves the same total capacity by operating two smaller installations rated 15 MW, which is below the threshold. Similar cases arise for process-regulated activities due to the capacity-based approach with sharp thresholds. If the capacity ratings of installations were known to us, we could identify the treatment effect in a regression-discontinuity design. However, no such data are publicly available for France (and, to the best of our knowledge, in any other European country). Nonetheless, we can take advantage of the fact that the participation rules induce variation in treatment status across plants and firms with similar levels of CO₂ emissions by using difference-in-differences approaches that have been successfully used in the evaluation of other cap-and-trade schemes (Fowlie et al., 2012). To internalize spillovers that may arise between regulated and unregulated plants that belong to the same firm, and to take advantage of a much larger set of firm-level outcome variables, we estimate average treatment effects on the treated at the firm level.

2.2 Difference-in-Differences Approach

Having longitudinal firm data allows us to estimate counterfactual emissions in the absence of the EU ETS and thereby tease apart the effect of the regulation. The firms in our sample can be in one of two states. Either they take part in carbon trading under the EU ETS or they prevail in a state of business as usual. Let $ETS_i = 1$ if firm i belongs to the treatment group as defined in the previous section and $ETS_i = 0$ otherwise. The potential outcomes $Y_{it}(1)$ and $Y_{it}(0)$, conditional on membership and non-membership respectively, denote the outcome variables of interest for installation i in the post-treatment period ($t = 1$) or the pre-treatment period ($t = 0$).

The difference-in-differences estimator calculates the difference in average emissions for unregulated firms before and after the EU ETS came into effect and subtracts this change from the difference in average emissions from regulated firms before and after the EU ETS

came into effect,

$$\begin{aligned}
\frac{DD}{ATT} &= E[Y_{i1}(1) - Y_{i1}(0) | ETS_i = 1] \\
&= fE[Y_{i1}(1) | ETS_i = 1] - E[Y_{i1}(1) | ETS_i = 0]g \\
&\quad - fE[Y_{i0}(0) | ETS_i = 1] - E[Y_{i0}(0) | ETS_i = 0]g
\end{aligned}$$

This approach assumes that regulated firms would have followed the same trend as unregulated firms in the absence of the treatment,

$$E[Y_{i1}(0) - Y_{i0}(0) | ETS_i = 1] = E[Y_{i1}(0) - Y_{i0}(0) | ETS_i = 0]:$$

In words, we assume that the trajectory of regulated firms would have continued to follow the trajectory of unregulated firms in the absence of the policy. We argue that the parallel trends assumption is plausible when evaluating the effects of the EU ETS at the firm-level within narrowly defined sectors. While not testable, we provide support for this assumption in the results below, showing that there are no significant differences in trends between regulated and unregulated firms, for observable outcomes, prior to the introduction of the EU ETS.

The combination of these two differences allows us to control for other changes over time – which would not be possible if we simply compared the difference in emissions for regulated firms before and after the introduction of the EU ETS –, as well as systematic differences between regulated and unregulated firms – which would not be possible if we simply compared the difference in emissions between regulated and unregulated firms after the EU ETS was introduced.

A priori, we cannot rule out spillovers between firms in the form of emissions leaking from ETS to control firms. This would violate the Stable Unit Treatment Assumption (SUTVA). However, such leakage would be associated with effects on other (economic) performance variables like employment or value added, which we do not see in the estimation results. Our results are thus unlikely to be confounded due to a violation of SUTVA.

In our baseline specification we estimate the following regression equation,

$$\Delta y_{j:s:t} = \sum_{s=1}^4 [\mathbb{1}_{ft} \beta \Theta g \quad ETS_j + \gamma_s] + \epsilon_{j:s:t} \quad (1)$$

where

$$\begin{aligned}
\Theta_1 &= f1996;:::;1999g && \text{(Pre-Announcement Period);} \\
\Theta_2 &= f2001;:::;2004g && \text{(Announcement Period);} \\
\Theta_3 &= f2005;:::;2007g && \text{(Trading Phase I), and} \\
\Theta_4 &= f2008;:::;2012g && \text{(Trading Phase II).}
\end{aligned}$$

The dependent variable $\Delta y_{j,s;t}$ denotes the difference in outcome y for firm j in sector s and year t relative to the year 2000, which preceded the announcement of the EU ETS. The treatment variable ETS_j takes a value of one for regulated firms and zero otherwise. The coefficient of interest captures the average effect of the EU ETS on regulated firms, relative to unregulated firms, across all years within the periods defined above. Sector-by-period fixed effects γ_s control for sector-specific time-varying unobservables, such as sector-wide changes in technology, policy or demand. Time-invariant firm-specific characteristics are controlled for by differencing the outcome variables relative to the year 2000. Standard errors are clustered at the firm-level.

2.3 Difference-in-Differences with Semi-Parametric Matching

We gauge the robustness of the difference-in-differences estimator by using a semi-parametric difference-in-differences approach, following Heckman et al. (1997, 1998):

$$\begin{aligned}
ATT^{matched} &= E[Y_{i1}(1) - Y_{i1}(0) | X_i; ETS_i = 1] \\
&= \frac{1}{N_1} \sum_{j \in I_1} \left\{ (Y_{j1}(1) - Y_{j0}(0)) - \sum_{k \in I_0} \lambda_{jk}(X_j; X_k) (Y_{k1}(0) - Y_{k0}(0)) \right\} \quad (2)
\end{aligned}$$

where I_1 denotes the set of ETS firms, I_0 the set of non-ETS firms, and N_1 the number of participating firms in the treatment group. The treated firms are indexed by j , the control firms are indexed by k . The weight placed on a non-ETS firm when constructing the counterfactual estimate for ETS firm j is λ_{jk} . These weights can be calculated using any matching approach. The rationale behind matching is to improve covariate balance and to increase common support between regulated and unregulated firms.

In our application, we implement this approach as a difference-in-differences regression on a matched sample obtained in a one-to-one nearest-neighbor matching. The regression

equation is given by

$$(Y_{j;t} - Y_{j;2000}) - (Y_{k;t} - Y_{k;2000}) = \sum_{g=1}^4 \mathbb{1}_{ft \geq \Theta_g} \alpha_{j;t} \quad (3)$$

where

$$\begin{aligned} \Theta_1 &= [1996; \dots; 1999] && \text{Pre-announcement period} \\ \Theta_2 &= [2001; \dots; 2004] && \text{Announcement period} \\ \Theta_3 &= [2005; 2006; 2007] && \text{Trading Phase I} \\ \Theta_4 &= [2008; \dots; 2012] && \text{Trading Phase II:} \end{aligned}$$

The left-hand side of equation (3) denotes the difference in outcome between treated firm j and matched control firm k in year t , relative to that difference in the base year 2000, i.e. just before the announcement of the EU ETS. The coefficients of interest are $\alpha_{j;t} = \frac{\text{matched}}{\text{ATT}}$, which provides the effect of the EU ETS on regulated firms in period t as compared to the matched control firms, and relative to the year 2000.

Matching Variables Following Fowlie et al. (2012), we match treated firms to control firms based on (i) the logarithm of their CO₂ emissions in the year 2000 (the year prior to the announcement of the EU ETS) and (ii) on the 2-digit NACE sector of the firm, which we re-define to reflect the fact that multi-plant firms may engage in multiple activities.⁹ Matching exactly on sector means that we control for sector specific shocks to the outcome variables that may have occurred after the introduction of the EU ETS. This set of matching variables is chosen to produce a comparison firm that has similar characteristics to a treatment firm while maximizing the number of successful matches. A parsimonious matching strategy such as this allows us to verify that covariates are balanced between treatment and control firms across both matched and unmatched variables. We do not match on pre-treatment trends. Instead, we let the data speak to the validity of the assumption that pre-treatment trends in the outcome variables are parallel. Appendix A provides further details on the matching process and quality, as well as alternative specifications and robustness checks.

⁹We define a new sector variable SUPERNACE at the firm level which is based on the combination of all plant-level activities. For example, if a firm owns two plants and both produce in NACE 12, then the SUPERNACE is 12 and the firm would be matched to a control firm in the same sector (with SUPERNACE 12). In contrast, for a firm with one plant producing in NACE 12 and another one in NACE 17, we define SUPERNACE to be 1217 and match it to a control firm within SUPERNACE 1217 (where the ordering of sectoral codes does not matter, e.g., SUPERNACE “1217” is equivalent to SUPERNACE “1712”).

Inference on Post-Matching Regression Coefficients It has been argued that matching can be seen as a pre-processing step to estimation and thus be ignored in the computation of standard errors (Ho et al., 2007). However, Abadie & Spiess (2019) demonstrate that bias in the estimation of the variance can occur if the covariates in the regression are correlated with the error term, conditional on the variables that have been matched on. They demonstrate that valid inference can be conducted if matching is done *without replacement* and standard errors are clustered at the level of the match.

Matching without replacement implies that a given control firm will only be used as a match in a given year for one particular treated firm. This has the potential downside of introducing bias in the asymptotic distribution of the post-matching regression estimator, especially when few suitable controls are available relative to the number of treated units. In our application, the number of control firms turns out to be low for a number of sectors. Consequently, matching without replacement reduces our analysis sample by 33% compared to matching with replacement, further reducing the comparability of the matched sample to the main analysis sample.

By contrast, matching with replacement allows for a larger sample size because multiple treated firms can be matched to the one control firm that best fulfills the matching criteria. Given the bias-variance trade-off we give priority to the former and use matching with replacement in our main specification. To address the point made by Abadie & Spiess (2019), we use a two-way cluster adjustment. The first cluster is at the level of the match (the firm) and also addresses serial correlation. The second cluster is at the control-firm-year level to account for correlation across observations that are matched to the same control observation. We propose that this additional adjustment addresses at least part of the concern associated with the effects of matching with replacement on inference. Notice that our adjustment collapses to the solution presented in Abadie & Spiess (2019) when each treatment firm is matched to a unique control firm. In this case the second cluster becomes redundant.

3 Data

This section details the different datasets used in our analysis. We compile a dataset of French manufacturing firms for each year between 1996 and 2012. This period covers several years prior to the announcement of the EU ETS, the announcement phase between 2001 and 2004, and Phases I and II. The data are obtained from various sources.¹⁰

¹⁰Firm- and plant-level data from the French Statistical Office used in this paper were provided for research purposes by authorization of the *Comité du Secret Statistique*, reference E598.

3.1 Energy and Emissions Data

We obtain detailed fuel use data from the Annual Survey of Industrial Energy Consumption (EACEI),¹¹ a survey conducted annually by the French National Institute of Statistics and Economic Studies (INSEE - Institut National de la Statistique et des Études Économiques). The survey provides quantities and values of energy consumed by fuel type¹² - broadly speaking, electricity, steam, fossil fuels and biofuels. Other variables available in the survey include the geographical location and sectoral classification of each establishment.

Having reliable data on CO₂ emissions is of central importance to our study. We calculate those emissions for both treated and untreated firms using the detailed energy consumption data from the EACEI in conjunction with standardized conversion factors provided by the French Environment & Energy Management Agency (ADEME - Agence de l'Environnement et de la Maîtrise de l'Énergie).¹³ Consequently, a firm will only be in our core dataset if it reports detailed energy consumption data under the EACEI. The sampling frame for the EACEI includes all French manufacturing establishments.¹⁴ The response rate is close to 90 percent. This speaks to the high representativeness of the dataset, but it is important to note that not all establishments are covered, and that sampling rules have changed over time. In 2000, the survey covers 88% of industrial emissions in France.

Until 2007, firms included in the EACEI survey were in sectors 12 to 37 according to the NAF rev.1 classification, equivalent at the two-digit level to the NACE rev.1.¹⁵ In the later years, different sampling weights were applied to draw around 12,000 establishments for the sample each year. The sample includes (i) all industrial establishments with 20 employees or more in the most energy consuming sectors;¹⁶ (ii) all establishments with more than ten employees in sector 20.11Z (manufacturing of industrial gases); (iii) all establishments with more than 250 employees on the 31st of December of that year; (iv) a random sample of establishments with employment between 20 and 249 employees in sectors that are not energy intensive. While the subsequent analysis is not based on the universe of French

¹¹In French: *Enquête annuelle sur les consommations d'énergie dans l'industrie*.

¹²Information for the following fuel types is requested from the surveyed firms: electricity (bought, auto-produced and resold), steam, natural gas, other types of gas available on the network, coal, lignite, coke, butane, propane, heavy fuel oil, heating oil, other petroleum products, the black liquor (a byproduct of the chemical decomposition of wood for making paper pulp), wood and its by-products, special renewable fuels, special non-renewable fuels.

¹³EU ETS participants in France are required to use the ADEME's conversion factors when reporting their emissions.

¹⁴The level of survey is the establishment rather than the enterprise given that energy consuming materials, electricity and gas meters and fuel tanks are held at that level.

¹⁵The following sectors are excluded: 15 - Manufacture of food products and beverages, 20.1A, 22.1 and 23

¹⁶23.32Z Manufacture of bricks, tiles and construction products, in baked clay; 23.51Z Manufacture of cement; 23.52Z Manufacture of lime and plaster

manufacturing firms, it draws on a database designed to provide a representative picture, especially of the most energy intensive firms in French manufacturing, while living up to the high standards of data collection for official statistics in France.

3.2 Financial Data

The employment and financial variables are obtained from French fiscal data. Tax returns filed by firms with the French Ministry for the Economy and Finance are collected in the annual fiscal census of manufacturing, mining and utilities firms. Until 2007, this census was called the Unified Corporate Statistics System (SUSE¹⁷) and the resulting dataset we exploit is the FICUS¹⁸ database which covers the years from 1994 to 2007. For the years from 2008 until 2012, the successor system is called ESANE (Elaboration des Statistiques Annuelles d'Entreprises) with the resulting dataset FARE (Fichier Approché des Résultats d'ESANE). These datasets provide general information about the firm (identifier, industry classification, head office address, total number of workers employed, age, etc.), the income statement (containing variables such as total turnover, total labor costs and value added) as well as balance sheet information (e.g. various measures of capital, debt and assets).¹⁹ As a measure of capital, we use the value of gross fixed tangible assets, which includes machinery, equipment and buildings.

3.3 Imports Data

Firm-level data on imports for the period of 1995 to 2012 are obtained from French Customs (DGDDI-Direction Générale des Douanes et Droits Indirects). The raw data are based on the customs declaration forms that firms are required to submit, and provide a comprehensive annual record of the value and quantity of exports and imports by destination, or origin, country at the eight-digit product (CN8) level. The customs dataset has been used previously in the trade literature (Eaton et al., 2011; Mayer et al., 2014). It includes the universe of trade flows from and to French firms, although reporting thresholds exist for compulsory declarations inside and outside the European Union. Outside the EU, exports or imports

¹⁷Systeme Unifié de Statistique d'Entreprises

¹⁸Fichier Complet Unifié de SUSE

¹⁹Observations displaying extreme changes in employment, value-added, emissions and emissions intensity are dropped. Taking the year 2000 as the base year, extreme change is defined as an increase or decrease of more than 300% from the years 1999 to 2000, or 2000 to 2001. For each prior or subsequent year, the median change is added to the threshold. For example extreme employment change from 2000 to 2003 would be defined as such if larger than 300% plus the median increase in employment from 2001 to 2002, plus the median increase in employment from 2002 to 2003. In addition, only observations with non-missing values for employment, value-added, emissions and capital are retained.

are only reported if their annual total is above €1,000 or 1,000 kg. Within the EU, these thresholds vary over time and by trade flow (imports vs. exports). To harmonize across different thresholds, we consider as non-traders firms that report total exports or imports within or outside the EU of less than €150,000. Firms that trade through intermediaries are also considered to be non-traders. Since all treated firms were importers in the reference year 2000, we drop untreated firms that do not import any goods in that year, to increase the comparability of regulated and unregulated firms.

3.4 Approximating the Carbon Intensity of Imports

To measure the carbon intensity of imports, we adopt the data and approach taken by the European Commission when establishing whether a sector is at risk of carbon leakage.²⁰ According to this approach, the carbon intensity of a sector is measured as the percentage share of carbon permit costs in value added. Carbon permit costs are calculated as the sum of indirect and direct carbon emissions multiplied by a fixed price of €30/tCO₂. This proxy for costs is then divided by the gross value added of a sector.

For each firm and year in our dataset, we use correspondence tables between NACE rev1.1 and CN8 product codes from Eurostat’s Reference and Management of Nomenclatures²¹ to obtain the value of imports of goods from a given sector. By multiplying these values with the sector’s carbon intensity and aggregating across sectors, this provides a carbon weighted measure of a firm’s imports value, reflecting the carbon intensity of its imports.

3.5 Environmental protection investments data

For a subset of firms, we obtain detailed data on investments for mitigating carbon emissions and air pollution. This dataset is also collected by INSEE as part of the Annual Survey on Environmental Protection Studies and Investments (Antipol)²². The sampling frame includes establishments from sections B, C and D of the NAF rev.2 classification, extending to some divisions of section E since 2012. Different sampling weights were applied to draw about 11,000 units. The response rate is above 80%.

The variables used here all relate to investment dedicated to reduce air pollution, including emissions of greenhouse gases. They are split between (a) investments made to “measure” air and GHG pollution, (b) “integrated” investments made in production processes and ma-

²⁰Cf. in the Commission Decision 2010/2/EU, pursuant to Directive 2003/87/EC of the European Parliament and of the Council, the list of sectors and subsectors at the NACE rev1.1 four-digit level which were deemed to be exposed to a significant risk of carbon leakage (2010) OJ L 1/10.

²¹This can be accessed on: <https://ec.europa.eu/eurostat/ramon>

²²In French: *Enquête sur les investissements et les dépenses courantes pour protéger l’environnement.*

chines that are less carbon- or air pollution-intensive than alternatives, and (c) “specific” investments made solely to limit and prevent air pollution and GHG emissions, e.g. a filter. All investments are reported in thousands of Euros. In estimating (b), the “integrated” investment, the respondent is asked to report the additional cost of an investment that is relevant for protecting the environment. For example, they would report the difference in the price of a new machine relative to that of an alternative that is more emissions-intensive. In addition, they report the share of total integrated environmental investments that are dedicated to air and climate pollution. Data about investments defined as (a) were collected since 1996. However, investments defined as (b) or (c) were only included in the survey from 2001. This means that for those two categories, we can only explore changes in investment relative to 2001. Given the frequent occurrence of zero values in the dataset, we apply an inverse hyperbolic sine (IHS) transformation rather than a logarithmic transformation:

$$\operatorname{arcsinh} y_{it} = \ln \left(y_{it} + \sqrt{y_{it}^2 + 1} \right)$$

The IHS transformation is approximately equal to $\log(2y_{it})$, except for very small values, and so can be interpreted in the same way as a logarithmic transformation. However, unlike the logarithmic transformation, the IHS of zero is well-defined.

3.6 EU Transaction Log Data

The European Union Transaction Log (EUTL) is the official registry of the EU ETS. It provides a list of all regulated installations, past and present.²³ A pollution right in the EU ETS is called a European Union Allowance (EUA). Each EU ETS installation has an “operator holding account” in its national registry, into which its own allowances are issued. Any individual or organization wishing to participate in the market is able to open their own “person holding account” in any of the registries. The internet portal of the EUTL makes publicly available contact details for each account, the number of allowances allocated under the “national allocation plan”, and the compliance position of each installation, which is calculated as the net balance of surrendered EUAs and verified emissions. This information is provided at the annual level. We combine it with the data described above to identify regulated firms.

²³When the EU ETS was established in 2005, each member state created its own national registry containing allowance accounts for each plant and other market participants. These registries interlinked with the Community Independent Transaction Log (CITL), operated by the Commission, which records and checks every transaction. Since 2012, the EU ETS registry has been operated in a centralized fashion as the EUTL.

3.7 Analysis Sample and Descriptive Statistics

The quality of the link between entities across datasets is an important determinant of the final sample in our empirical analysis. Linking the EACEI, FICUS/FARE, trade data and Antipol is straightforward as all four datasets use the SIREN (*Systeme d'Identification du Répertoire des Entreprises*) number as their identifier²⁴. By linking the EACEI to the cleaned FICUS/FARE, some firms are dropped and the resulting sample emits 55 million tonnes of CO₂ in 2000, which represents 51% of aggregate industrial emissions in France. Not all firms from our main dataset are surveyed in Antipol. While the business dataset is maintained by INSEE, the French national registry of the CITL/EUTL is managed by Caisse des Dépôts. The latter institution provides a link between the permit identifier (GIDIC) from the national registry and the SIREN identifier from INSEE, allowing for the linking of the EUTL data to the business data. We drop ETS firms that are not observed either before the policy was announced in 2000 or after it was introduced in 2005 as this would prevent us from running our analysis.

The resulting sample used for the analysis is an unbalanced panel of 3,837 firms, 163 of which are part of the EU ETS. Only when analyzing investment outcomes taken from the Antipol survey do we resort to a smaller sample of 2,936 firms that are sampled at least in 2001 and one other year, corresponding to 41% of the firm-year observations in the full sample. The main variables are summarized in Table A1.

Appendix Table B1 reports the variables evaluated in the base year 2000. The first two columns report means and standard deviations separately by treatment status, and the third column shows the differences in outcomes between ETS and non-ETS firms. We can see that ETS firms are on average larger in terms of employment, value added, capital and imports. They also emit more CO₂ emissions and are more carbon intensive. The fourth column shows that nearest-neighbor matching with replacement reduces these differences. Figures B1 through B3 further illustrate the improvement in balance between treatment and control firms, plotting the full distribution of our key outcome variables of interest in the full unmatched sample (left figures) and matched sample (right figures). Across all variables we see that the balance between regulated and unregulated firms improves substantially in the matched sample.

Our matching algorithm does not completely eliminate differences between treated and control firms in the sample. However, it is important to note that perfect balance is not a necessary condition for identification in this context. The critical identification assumption

²⁴Although each plant in the EACEI and Antipol is identified by a SIRET (*Système d'Identification du Répertoire des Etablissements*) number, the SIREN number corresponds to the first nine digits of the SIRET number.

is that treated firms would have followed the same trajectory as control firms in the absence of the EU ETS – the parallel trends assumption. In support of this assumption, we do not observe differential trends between the treated and control firms prior to the announcement of the EU ETS.

4 Results

4.1 Main Outcomes

Table 1 presents the estimation results for equation (1) and establishes the following main results. First, regulated firms reduced emissions by 8.2% ($p < 0.05$) compared to unregulated firms during Trading Phase II. We fail to reject the null hypothesis that firms did not reduce emissions during the Announcement Phase and Trading Phase I. Second, we fail to reject the null hypothesis that the policy had no impact on the economic performance of firms, as measured by value added or number of employees.²⁵ Consistent with this, we estimate a 10.7% ($p < 0.05$) reduction in the emissions intensity of value added in Trading Phase II. Third, with somewhat lower confidence we estimate that regulated firms increased their capital stock by 6.4% ($p < 0.1$) relative to unregulated firms in Trading Phase II.

A key assumption required for us to interpret these effects as causal is that regulated firms would have followed the same trajectory as unregulated firms in the absence of the policy. Consistent with this assumption, we do not estimate any differential trend in the pre-announcement period for any of our outcomes. Figure 2 presents graphical support for this assumption.

In Table 2, we test the robustness of our main findings using a matched sample. The estimates in our matched sample are broadly similar to those in our main specification, however, they are more noisily estimated, consistent with the bias-variance trade-off associated with matching. We estimate an 11.7% reduction in emissions during the second trading phase, no effects on value added or employment, and an 11.8% reduction in the carbon intensity of value added. The estimated effect on capital is no longer statistically significant in this specification, but the point estimate is of a similar magnitude.

One qualitatively different finding compared to the baseline results is that CO₂ emissions were 6% higher among treated firms during the announcement period. This increase could be motivated by rent seeking. If regulated firms correctly anticipated that they were going to receive free permits in proportion to historical emissions under a grandfathering allocation

²⁵The absence of effects on economic performance helps to mitigate the concern that the estimated effects might be confounded by differential reactions to the 2008 recession between treated and control groups. If this was the case we would expect to see reductions in economic outcomes as well as environmental outcomes.

scheme, they could generate permit rents by ramping up emissions during the announcement phase.²⁶ While economically plausible, this result should be taken with a grain of salt as it is only supported by the matching specification.

In Appendix B we show that our main results are robust to using alternative matching specifications (Table B2) and to imposing caliper restrictions on the matching distance between treatment and control firms (Table B3).

4.2 Leakage

According to our results, the EU ETS induced regulated firms to reduce emissions with no detectable effects on economic performance. This begs the question of whether part, or all, of the estimated emissions reductions arise due to carbon leakage, i.e. the reallocation of emissions to unregulated firms. To assess the efficacy of the EU ETS as a climate policy instrument, it is important to know whether the CO₂ abatement we have estimated represents a global reduction in emissions.

Carbon leakage could occur through multiple channels. Three of them are particularly relevant in the context of our study. The first channel is via the supply chain. That is, a regulated firm might reduce the emissions intensity of its operations by out-sourcing more intermediate products from unregulated firms. Such a strategy could save on compliance costs, particularly if applied to the most carbon-intensive steps of the value chain. But it would inevitably reduce the firm’s value added, defined as “revenue minus material inputs”, where material inputs are sourced both domestically and through international trade. We do not estimate any reduction in value added. Moreover, regression results reported in Table 3 show that the EU ETS did not affect the value, nor carbon intensity, of intermediate inputs imported by regulated firms. These findings also hold for the matched sample. Hence, the notion that out-sourcing helped firms to reduce their emissions is not supported by the data.

The second channel of carbon leakage is via the product market. Because carbon pricing increases production costs at regulated firms, market forces might shift production to unregulated firms within France or abroad. If this process was driving the negative effect we estimate for emissions, we would expect to see a negative effect of the EU ETS on at least one of the economic variables such as value added, employment or investment. However, instead we estimate insignificant effects on employment and value added, and positive effects on capital investment. Apart from mitigating concerns about leakage, this result is

²⁶Under the French national allocation plan for trading phase I, free permits were granted to industrial installations in proportion emissions during the period from 1996 until 2002. For new installations, data from 2004 and 2005 would also be taken into account (Ministère de l’environnement, 2005). Bushnell et al. (2013) present evidence that free permit endowments were highly valued by investors of publicly traded firms that were regulated under the ETS.

useful as a an indirect test of whether treatment spillovers, which could pose a threat to our identification strategy, are empirically relevant in this context. Product-market leakage is isomorphic to a treatment spillover between regulated and unregulated firms which reallocates market share from regulated to unregulated firms. This would violate SUTVA and lead to an overstatement of the treatment effect as emissions fall at regulated firms and increase at unregulated firms, in lock-step with production. Yet again, the same effect should be observed for other variables relating to the scale of production. We find no evidence that this is the case. We only estimate reductions in emissions.

A third possible channel of leakage arises if firms operating multiple facilities reallocate production from regulated to unregulated ones. We internalize within-firm spillovers by estimating the effects of the EU ETS at the firm-level. Consequently, within-firm leakage cannot explain estimated emissions reductions at the firm-level. Our estimates are net of any within-firm leakage.²⁷

4.3 Abatement channels

The absence of evidence on carbon leakage, combined with the estimated reduction in the carbon intensity of value added, supports the view that emission reductions arose from improvements to the emissions intensity of production. Such improvements can be achieved by switching to less polluting fuels or by investing in technology that is more efficient (or indeed from investments in technology that allows fuel switching). Our data allow us to explore these different channels of abatement.

Table 4 reports results relevant for investigating whether emissions reductions (reported in column 1 for reference) were achieved by substituting low-carbon fuels for high-carbon fuels.²⁸ Contrary to this mechanism, we estimate significant reductions in emissions from natural gas (column 2), a low-carbon fuel, particularly when considering the intensity of value added with respect to this variable (column 3). To understand why regulated firms cut emissions from a low-carbon fuel, note that natural gas is accounting for most (79%) of CO₂ emissions at baseline and hence the potential for intensive-margin switching into gas is very limited.²⁹ To check for possible fuel switching on the extensive margin, we also report results for the share of natural gas in total CO₂ emissions (column 4). Since we do not estimate an increase in this variable, we conclude that other fossil fuels played no major

²⁷In Table C1 we present results on the effects of the EU ETS on plant-level emissions. The results remain robust.

²⁸Table 5 is analogous to Table 4 and reports results based on the matched sample. Unless we state otherwise in the subsequent discussion, the results are robust across specifications.

²⁹This share refers to the 148 treated firms that used natural gas in 2000. The unconditional emissions share of natural gas is 72% at baseline.

role in explaining emissions reductions in Phase II.³⁰ Another possible fuel-switching channel is that regulated firms used more electric energy in the production process. The principal mechanism for this is by procuring more electricity from the grid.³¹ We look at this outcome and find no increase in the amount of electricity bought by firms (column 7). In sum, the results indicate that fuel switching to natural gas or electricity cannot explain the estimated CO₂ abatement at regulated firms. An implication for climate change mitigation is that CO₂ abatement by regulated manufacturing firms did not lead to increased emissions in the electricity sector.³²

This leaves technology adoption as a possible mechanism behind the reductions in carbon emissions and emissions intensity of regulated firms. The positive treatment effect on capital stock is suggestive, but not conclusive, evidence that regulated firms invested in reducing the emissions intensity of production.

Table 6 provides further evidence in support of this hypothesis based on data pollution control investment. Specifically, we estimate that regulated firms significantly increased their investments in integrated production technologies that reduce air and climate change related pollution emissions, such as more efficient boilers, during Trading Phases I and II (Column 2). We do not estimate any differential impacts on investments into the measurement of emissions (not needed for CO₂ given the ease of input-based accounting) or investments into specific, ‘end-of-pipe’ measures to reduce emissions (not yet available for CO₂ at a commercial scale). We find qualitatively consistent results in the matched sample, although the estimated effect of the EU ETS in Phase II is statistically insignificant. A caveat with this analysis is that, unfortunately, data for integrated and specific investments were only collected from 2001 onward. Consequently, we are unable to investigate whether trends in those outcomes are parallel during the pre-announcement period. Data for this period is available though for the outcome variable reported in column 1, measurement investment, where we do not find any differential effect in the pre-announcement phase.

In sum, our findings suggest that the principal mechanism underlying the estimated emissions reductions is that treated firms reduced the carbon intensity of production by upgrading their capital stock.

³⁰If they did, the gas share in total emissions should increase, but our point estimates for the post-announcement periods are negative. The coefficient in Phase I is statistically significant at 10%, but this is not robust when we repeat the estimation for the matched sample in Table 5.

³¹Firms could also generate more electricity on site, but this is quite rare among the firms in our sample and would lead to higher direct emissions, contrary to what we find.

³²It is likely that buying electricity would not lead to an increase in global emissions because 79% of the electricity generated in France in 2012 was carbon neutral, with the balance likely to have been produced by an ETS regulated power plant.

4.4 Permit Allocations

The way pollution rights are allocated to emitters is an important design feature of any cap-and-trade scheme. During our study period, manufacturing firms received permits free-of-charge in proportion to their historical emissions. In theory, the initial allocation of permits should not distort a firm’s marginal abatement decisions (Montgomery, 1972; Coase, 1960; Dales, 1968; Fowlie & Perloff, 2013). In practice, however, permit allocations could matter because of market power, credit constraints and other frictions, or due to behavioral factors (Hahn, 1984; Kahneman et al., 1990).

We investigate whether treatment effects vary with respect to a firm’s initial permit surplus, defined as the ratio of allocated permits to actual emissions at the beginning of the policy. In Appendix D we document substantial variation in this variable. We divide firms into two groups – those that are above and below the median surplus – and estimate separate treatment effects for each group. We find robust evidence that the estimated effects on CO₂ emissions, carbon intensity of production, and capital are driven by those firms that were relatively short of permits (Table D1). Since permits were not allocated at random to individual firms, we cannot interpret this treatment heterogeneity as causal. Nevertheless, the result suggests that the initial allocation of permits mattered in the context of the EU ETS. Firms that received fewer permits than needed likely perceived the policy as more stringent than those with a more generous permit surplus. The treatment heterogeneity we uncover is consistent with permit-constrained firms abating emissions so as to ensure compliance with the policy. In contrast, firms with a permit surplus did not have to reduce emissions in order to comply – and chose not to.

5 Aggregate Carbon Savings

We combine our estimates with data on aggregate CO₂ emissions to gauge the potential contribution of the EU ETS in driving aggregate emission reductions since 2005. Details on the calculations made below can be found in Appendix E.

The black line in Figure 3 depicts the observed trajectory of aggregate industrial CO₂ emissions in France between 1996 and 2012. We observe that emissions have been falling over time, and that the decline has been steeper in recent years.

Compared to emissions in 2004, there is little change in emissions during the first trading phase of the EU ETS. This is what our difference-in-differences estimator implies, failing to reject the null hypothesis that the ETS had no effect on emissions during this period (2005-2007). However, emissions declined substantially in the second phase of the EU ETS. Our

difference-in-differences estimates suggest that aggregate industrial emissions would have been, on average, 4.74 million tonnes higher each year in the absence of the EU ETS. As such we posit that the EU ETS accounted for roughly 26% of the aggregate reductions in industrial CO₂ emissions during this period.

We also reflect on what our estimates could imply for the effects of the ETS on aggregate industrial emissions in the EU more broadly. Such an exercise assumes that our estimates have external validity, i.e., they are representative of how regulated firms in other member countries responded. This is a strong assumption. While all firms face the same carbon price, differences in permit allocations, market structure, and regulatory oversight may have resulted in smaller, or larger, treatment effects outside of France. With caveats noted, we calculate the implied effect of the EU ETS on aggregate industrial emissions for the EU market as a whole. Our calculations suggest that aggregate industrial emissions would have been 36 million tonnes higher each year between 2008 and 2012, in the absence of the EU ETS.

These emissions reductions occurred in spite of carbon prices averaging at a rather low \$21.35 per tonne (\$2017) during Phase II. Arguably, the average abatement costs per tonne of CO₂ must have been lower, for otherwise it would have been more profitable for firms to purchase permits instead of reducing emissions.

Does that make the EU ETS an expensive policy? Previous research on air pollution regulation has established that the overall cost of market-based instruments compares favorably with that of non-market based approaches [Carlson et al. \(2000\)](#); [Fowle et al. \(2012\)](#); [Gillingham & Stock \(2018\)](#). In Figure 4 we compare the estimated cost per tonne of CO₂ (\$2017) for 25 climate change mitigation policies. The estimate for the EU ETS is based on the maximum price during Phase II – \$52.68. This is the most conservative cost as above this cost it would have been cheaper for firms to buy emission permits instead. Estimates for other climate change mitigation policies come from [Gillingham & Stock \(2018\)](#). Even when we use the maximum cost per tonne of CO₂, the EU ETS is ranked 7th. If we use the average Phase II price instead (\$21.35), which is still likely to be very conservative the EU ETS is ranked 5th.

6 Conclusion

In the context of the world’s largest carbon market, we have presented evidence that market-based regulatory instruments have the potential to reduce carbon emissions without imposing significant economic losses on regulated firms. These firms invested in new technology and thereby lowered the carbon intensity of production. We find no evidence of carbon leakage,

suggesting that the EU ETS helped to mitigate global climate change.

Our results contrast with the impacts of command-and-control regulations that impose one-size-fits-all regulatory standards for industrial air pollution emissions. While also delivering improvements in environmental quality, such non-market-based policies have been shown to have negative effects on firm performance (Becker & Henderson, 2000; Greenstone, 2002; Greenstone et al., 2012; Walker, 2013; He et al., 2020).

We note caveats. Despite the significant effect that the EU ETS has had on emissions, these results should not be taken as a blank endorsement of market-based regulatory instruments. Our findings have focused on the response of manufacturing firms in one market, and on one market-based regulatory instrument – emission trading schemes. Our context is one in which compliance is high and corruption low. Understanding the degree to which monitoring and enforcement constraints may affect the efficacy of environmental policies in other contexts presents interesting and important opportunities for future research. Further, our results do not imply that the emission reductions have been without cost. We lack suitable data to investigate whether such costs have reduced profits or whether firms were able to pass them on to their customers in the form of higher product prices. It is also possible that the ETS regulation made firms aware of cost-saving opportunities that they had previously not paid attention to, improving energy efficiency and reducing the emissions intensity of output in the process. Finally, our results do not guarantee that the ETS operates efficiently. Credit constraints, information asymmetries, market power in product markets, transaction costs, and other sources of market failure could all affect the efficiency of the scheme.

References

- Abadie, A. & Spiess, J. (2019). Robust post-matching inference. *Mimeo*.
- Andersen, B. & Di Maria, C. (2011). Abatement and allocation in the pilot phase of the EU ETS. *Environmental and Resource Economics*, 48(1), 83–103.
- Baumol, W. (1972). On taxation and the control of externalities. *American Economic Review*, 62(3), 307–322.
- Baumol, W. & Oates, W. (1971). The use of standards and prices for protection of the environment. *Swedish Journal of Economics*, 73, 42–54.
- Becker, R. & Henderson, J. (2000). Effects of air quality regulations on polluting industries. *Journal of Political Economy*, 108(2), 379–421.

- Borghesi, S., Franco, C., & Marin, G. (2020). Outward Foreign Direct Investment Patterns of Italian Firms in the European Union's Emission Trading Scheme. *The Scandinavian Journal of Economics*, 122(1), 219–256.
- Burke, M., Craxton, M., Kolstad, C., Onda, C., Allcott, H., Baker, E., Barrage, L., Carson, R., Gillingham, K., Graff-Zivin, J., Greenstone, M., Hallegatte, S., Hanemann, W., Heal, G., Hsiang, S., Jones, B., Kelly, D., Kopp, R., Kotchen, M., Mendelsohn, R., Meng, K., Metcalf, G., Moreno-Cruz, J., Pindyck, R., Rose, S., Rudik, I., Stock, J., & Tol, R. (2016). Opportunities for advances in climate change economics. *Science*, 352(6283), 292–293.
- Bushnell, J., Chong, H., & Mansur, E. (2013). Profiting from Regulation: Evidence from the European Carbon Market. *American Economic Journal: Economic Policy*, 5(4), 78–106.
- Calel, R. (2020). Adopt or innovate: Understanding technological choices under cap-and-trade. *American Economic Journal: Economic Policy*, 12(3), 170–201.
- Calel, R. & Dechezleprêtre, A. (2016). Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market. *Review of Economics and Statistics*, 98(1), 173–191.
- Carlson, C., Burtraw, D., Cropper, M., & Palmer, K. (2000). Sulfur dioxide control by electric utilities: What are the gains from trade? *Journal of Political Economy*, 108(6), 1292–1326.
- Coase, R. (1960). The problem of social cost. *The Journal of Law and Economics*, 3, 1–44.
- Dales, J. (1968). Pollution, property, and prices: an essay in policy-making and economics. *Mimeo*.
- Dechezleprêtre, A., Nachtigall, D., & Venmans, F. (2018). *The joint impact of the European Union Emissions Trading System on carbon emissions and economic performance*. OECD Economics Department Working Papers 1515, OECD Publishing.
- Eaton, J., Kortum, S., & Kramarz, F. (2011). An anatomy of international trade: Evidence from French firms. *Econometrica*, 79(5), 1453–1498.
- Egenhofer, C., Alessi, M., Fujiwara, N., & Georgiev, A. (2011). The EU emissions trading system and climate policy towards 2050: Real incentives to reduce emissions and drive innovation? *Centre for European Policy Studies Special Report*.

- Ellerman, A. D. & Buchner, B. K. (2008). Over-Allocation or Abatement? A Preliminary Analysis of the EU ETS Based on the 2005–06 Emissions Data. *Environmental and Resource Economics*, 41(2), 267–287.
- Ellerman, A. D., Convery, F. J., & De Perthuis, C. (2010). *Pricing Carbon: The European Union Emissions Trading Scheme*. Cambridge University Press.
- Ellerman, A. D. & Feilhauer, S. M. (2008). A top-down and bottom-up look at emissions abatement in germany in response to the eu ets. *Mimeo*.
- Ellerman, D., Marcantonini, C., & Zaklan, A. (2016). The European Union Emissions Trading System: Ten Years and Counting. *Review of Environmental Economics and Policy*, 10(1), 89–107.
- European Commission (2000). Green Paper on Greenhouse Gas Emissions Trading within the European Union. *Green Paper*.
- European Commission (2001). Green Paper on Greenhouse Gas Emissions Trading within the European Union - Summary of Contents. *Green Paper*.
- Fabra, N. & Reguant, M. (2014). Pass-Through of Emissions Costs in Electricity Markets. *American Economic Review*, 104(9), 2872–2899.
- Fowlie, M., Holland, S., & Mansur, E. (2012). What Do Emissions Markets Deliver and to Whom? Evidence from Southern California’s NOx Trading Program. *American Economic Review*, 102(2), 965–993.
- Fowlie, M. & Perloff, J. (2013). Distributing pollution rights in cap-and-trade programs: Are outcomes independent of allocation? *Review of Economic and Statistics*, 95(5), 1640–1652.
- Fowlie, M., Reguant, M., & Ryan, S. (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*, 124(1), 249–302.
- Gillingham, K. & Stock, J. (2018). The cost of reducing greenhouse gas emissions. *Journal of Economic Perspectives*, 32(4), 53–72.
- Greenstone, M. (2002). The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the Census of Manufactures. *Journal of Political Economy*, 110(6), 1175–1219.
- Greenstone, M., List, J., & Syverson, C. (2012). The Effects of Environmental Regulation on the Competitiveness of US Manufacturing. *NBER Working Paper No. 18392*.

- Hahn, R. (1984). Market power and transferable property rights. *The Quarterly Journal of Economics*, 99, 753–765.
- Hahn, R. W. (1989). Economic prescriptions for environmental problems: How the patient followed the doctor’s orders. *Journal of Economic Perspectives*, 3(2), 95–114.
- He, G., Wang, S., & Zhang, B. (2020). Watering Down Environmental Regulation in China. *Quarterly Journal of Economics*, 135(4), 2135–2185.
- Heckman, J., Ichimura, H., & Todd, P. (1997). Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme. *The Review of Economic Studies*, 64, 605–654.
- Heckman, J., Ichimura, H., & Todd, P. (1998). Matching as an econometric evaluation estimator. *Review of Economic Studies*, 65, 261–294.
- Hendren, N. & Finkelstein, A. (forthcoming). Welfare analysis meets causal inference. *Journal of Economic Perspectives*.
- Hendren, N. & Sprung-Keyser, B. (2020). A unified welfare analysis of government policies. *Quarterly Journal of Economics*, 135(3), 1209–1318.
- Ho, D., Imai, K., King, G., & Stuart, E. (2007). Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political Analysis*, 15(3), 199–236.
- Jaraite, J. & Di Maria, C. (2016). Did the EU ETS make a difference? An empirical assessment using Lithuanian firm-level data. *Energy Journal*, 37(1), 1–23.
- Kahneman, D., Knetsch, J., & Thaler, R. (1990). Experimental tests of the endowment effect and the coase theorem. *Journal of Political Economy*, 98(6), 1325–1348.
- Klemetsen, M. E., Rosendahl, K. E., & Jakobsen, A. L. (2020). The impacts of the EU ETS on Norwegian plants’ environmental and economic performance. *Climate Change Economics*, 11(1).
- Koch, N. & Basse Mama, H. (2019). Does the EU Emissions Trading System induce investment leakage? Evidence from German multinational firms. *Energy Economics*, 81, 479 – 492.

- Marin, G., Marino, M., & Pellegrin, C. (2018). The impact of the european emission trading scheme on multiple measures of economic performance. *Environmental and Resource Economics*, 71, 551–582.
- Martin, R., de Preux, L., & Wagner, U. (2014a). The impact of a carbon tax on manufacturing: Evidence from microdata. *Journal of Public Economics*, 117, 1–14.
- Martin, R., Muûls, M., de Preux, L., & Wagner, U. (2014b). Industry compensation under relocation risk: A firm-level analysis of the EU Emissions Trading Scheme. *American Economic Review*, 104(8), 1–24.
- Martin, R., Muûls, M., & 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.
- Mayer, T., Melitz, M., & Ottaviano, G. (2014). Market size, competition, and the product mix of exporters. *American Economic Review*, 104(2), 495–536.
- Ministère de l’environnement (2005). Décret n° 2004-832 du 19 août 2004 pris pour l’application des articles L. 229-5 à L. 229-19 du code de l’environnement et relatif au système d’échange de quotas d’émission de gaz à effet de serre. *Journal O ciel de la République Française*.
- Montgomery, W. (1972). Markets in licenses and efficient pollution control programs. *Journal of Economic Theory*, 5(3), 395–418.
- Nordhaus, W. (1977). Economic growth and climate: The carbon dioxide problem. *American Economic Review*, 67(1), 341–346.
- Nordhaus, W. (2001). Global warming economics. *Science*, 294(5545), 1283–1284.
- Petrick, S. & Wagner, U. J. (2014). The Impact of Carbon Trading on Industry: Evidence from German Manufacturing Firms. *Kiel Working Papers 1912*.
- Pigou, A. (1920). *The Economics of Welfare*. London: Macmillan.
- Ryan, S. P. (2012). The Costs of Environmental Regulation in a Concentrated Industry. *Econometrica*, 80(3), 1019–1061.
- Tietenberg, T. (1973). Controlling pollution by price and standard systems: A general equilibrium analysis. *Swedish Journal of Economics*, 75, 193–203.

Walker, W. (2013). The Transitional Costs of Sectoral Reallocation: Evidence from the Clean Air Act and the Workforce. *Quarterly Journal of Economics*, 128(4), 1787–1835.

Zaklan, A. (2020). Coase and Cap-and-Trade: Evidence on the Independence Property from the European Electricity Sector. *DIW Discussion Paper 1850*.

Tables and Figures

Table 1: The Effect of the EU ETS on the Environmental and Economic Performance of Firms

	(1)	(2)	(3)	(4)	(5)
	$\Delta \log(\text{CO}_2)$	$\Delta \log(\text{Value Added})$	$\Delta \log(\text{Emp.})$	$\Delta \log(\text{Capital})$	$\Delta \log(\text{CO}_2/\text{VA})$
Pre-Announcement	-0.019 (0.020)	0.015 (0.023)	0.025 (0.016)	0.010 (0.017)	-0.034 (0.028)
Announcement Period	-0.035 (0.022)	0.003 (0.024)	-0.009 (0.014)	-0.001 (0.014)	-0.038 (0.028)
Trading Phase I	-0.044 (0.032)	-0.041 (0.036)	-0.015 (0.025)	0.024 (0.028)	-0.003 (0.038)
Trading Phase II	-0.082 (0.041)	0.025 (0.048)	0.024 (0.034)	0.064 (0.037)	-0.107 (0.048)
Mean in 2000	135.354	91.934	1,213	197.954	0.003
Observations	42,733	42,733	42,733	42,733	42,733
Total # of Firms	3,837	3,837	3,837	3,837	3,837
# of Regulated Firms	163	163	163	163	163
Adjusted R ²	0.028	0.075	0.099	0.449	0.049

NOTES: These estimates are the result of OLS regressions. They provide the difference between regulated firm and unregulated firm outcomes prior to the announcement of the EU ETS, during the announcement period and during Phase I and Phase II of the EU ETS. Each coefficient represents the difference relative to the year 2000. Standard errors are clustered at the firm-level. Means are reported for ETS firms in 2000. Units: CO₂ { thousands of tons of CO₂; Value Added { millions of Euros; Employment { full-time equivalent employees; Capital { millions of Euros; CO₂/VA units { thousands of tonnes of CO₂ per thousand Euros of value added. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01.

Table 2: The Effect of the EU ETS on Firm Outcomes (Matched Sample)

	(1)	(2)	(3)	(4)	(5)
	$\Delta \log(\text{CO}_2)$	$\Delta \log(\text{Value Added})$	$\Delta \log(\text{Emp.})$	$\Delta \log(\text{Capital})$	$\Delta \log(\text{CO}_2/\text{VA})$
Pre-Announcement	0.016 (0.022)	0.037 (0.026)	-0.012 (0.018)	0.009 (0.021)	-0.022 (0.033)
Announcement Period	0.060 (0.022)	0.032 (0.044)	0.018 (0.019)	0.022 (0.017)	0.027 (0.040)
Trading Phase I	-0.002 (0.054)	-0.028 (0.056)	-0.040 (0.041)	0.019 (0.041)	0.026 (0.059)
Trading Phase II	-0.117 (0.054)	0.001 (0.083)	-0.064 (0.049)	0.062 (0.069)	-0.118 (0.072)
Mean in 2000	90.440	51.751	706	120.957	0.003
Observations	1,954	1,954	1,954	1,954	1,954
Total # of Firms	323	323	323	323	323
# of Regulated Firms	149	149	149	149	149
Adjusted R ²	0.026	0.002	0.011	0.007	0.011

Notes: These estimates are the result of OLS regressions, estimated on a matched sample. They provide the difference between regulated firm and unregulated firm outcomes prior to the announcement of the EU ETS, during the announcement period and during Phase I and Phase II of the EU ETS. Each coefficient represents the difference relative to the year 2000. Standard errors are two-way clustered by firm and matching group. Means are reported for ETS firms in 2000. Units: CO₂ { thousands of tonnes of CO₂; Value Added { millions of Euros; Employment { full-time equivalent employees; Capital { millions of Euros; CO₂/VA units { thousands of tonnes of CO₂ per thousand Euros of value added. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01.

Table 3: The Effect of the EU ETS on Imports

	(1)	(2)	(3)	(4)
	$\Delta \log(\text{Imports})$ Total	$\Delta \log(\text{Imports})$ CO ₂ intensive	$\Delta \log(\text{Imports})$ Total	$\Delta \log(\text{Imports})$ CO ₂ intensive
	Full Sample		Matched Sample	
Pre-Announcement	-0.008 (0.083)	0.049 (0.094)	-0.144 (0.116)	-0.150 (0.138)
Announcement Period	-0.062 (0.063)	-0.090 (0.087)	0.120 (0.099)	0.161 (0.133)
Trading Phase I	-0.109 (0.103)	-0.102 (0.101)	0.167 (0.140)	0.128 (0.182)
Trading Phase II	-0.047 (0.122)	-0.151 (0.134)	-0.123 (0.199)	-0.154 (0.175)
Mean in 2000	113.041	0.884	50.094	0.563
Observations	42,106	41,497	1,939	1,901
Total Firms	3,837	3,782	322	316
Treated Firms	163	163	149	148
Adjusted R ²	0.095	0.153	0.007	0.006

Notes: These estimates are the result of OLS regressions. They provide the difference between regulated firm and unregulated firm outcomes prior to the announcement of the EU ETS, during the announcement period and during Phase I and Phase II of the EU ETS. Each coefficient represents the difference relative to the year 2000. Standard errors are clustered by firm. Means are reported for ETS firms in 2000. Units: total imports and CO₂ Intensive Imports are measured in millions of Euros. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01.

Table 4: Effects of the EU ETS on Firm Energy Outcomes

	(1)	(2)	(3)	(4)	(5)
	$\Delta \log(\text{CO}_2$ Total)	$\Delta \log(\text{CO}_2$ from Gas)	$\Delta \log(\text{CO}_2$ Gas/V.A.)	Δ Share Gas in CO_2	$\Delta \log(\text{Electricity}$ Bought)
Pre-Announcement	-0.019 (0.020)	-0.025 (0.038)	-0.049 (0.046)	0.003 (0.013)	-0.017 (0.019)
Announcement Period	-0.035 (0.022)	-0.029 (0.039)	-0.048 (0.044)	-0.009 (0.011)	-0.014 (0.023)
Trading Phase I	-0.044 (0.032)	-0.065 (0.052)	-0.036 (0.055)	-0.035 (0.020)	-0.050 (0.036)
Trading Phase II	-0.082 (0.041)	-0.120 (0.064)	-0.152 (0.066)	-0.039 (0.030)	-0.029 (0.050)
Mean in 2000	135.354	60.227	0.003	0.721	189.732
Observations	42,733	32,103	32,103	42,733	42,728
Total # of Firms	3,837	2,888	2,888	3,837	3,836
# of Regulated Firms	163	148	148	163	163
Adjusted R ²	0.028	0.100	0.076	0.085	0.080

Notes: These estimates are the result of OLS regressions. They provide the difference between regulated firm and unregulated firm outcomes prior to the announcement of the EU ETS, during the announcement period and during Phase I and Phase II of the EU ETS. Each coefficient represents the difference relative to the year 2000. Note that one firm in the sample does not buy any electricity and is therefore not included in the last column regression. Standard errors are clustered by firm. Means are reported for ETS firms in 2000. Units: total emissions, and emissions from gas are measured in thousands of tonnes of CO₂, Value Added in millions of Euros, electricity bought in GWh. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01.

Table 5: Effects of the EU ETS on Firm Energy Outcomes (Matched Sample)

	(1)	(2)	(3)	(4)	(5)
	$\Delta \log(\text{CO}_2$ Total)	$\Delta \log(\text{CO}_2$ from Gas)	$\Delta \log(\text{CO}_2$ Gas/V.A.)	Δ Share Gas in CO_2	$\Delta \log(\text{Electricity}$ Bought)
Pre-Announcement	0.016 (0.022)	0.058 (0.044)	0.019 (0.055)	0.052 (0.030)	-0.072 (0.039)
Announcement Period	0.060 (0.022)	0.061 (0.036)	0.026 (0.059)	0.003 (0.009)	0.040 (0.025)
Trading Phase I	-0.002 (0.054)	-0.108 (0.081)	-0.088 (0.079)	-0.002 (0.028)	0.006 (0.043)
Trading Phase II	-0.117 (0.054)	-0.169 (0.080)	-0.162 (0.090)	0.019 (0.036)	-0.041 (0.053)
Mean in 2000	90.440	53.662	0.002	0.721	109.513
Observations	1,954	1,570	1,570	1,954	1,954
Total # of Firms	323	277	277	323	323
# of Regulated Firms	149	133	133	149	149
Adjusted R^2	0.026	0.026	0.012	0.014	0.013

Notes: These estimates are the result of OLS regressions, estimated on the matched sample. They provide the difference between regulated firm and unregulated firm outcomes prior to the announcement of the EU ETS, during the announcement period and during Phase I and Phase II of the EU ETS. Each coefficient represents the difference relative to the year 2000. Standard errors are clustered by firm. Means are reported for ETS firms in 2000. Units: total emissions, and emissions from gas are measured in thousands of tonnes of CO_2 , Value Added in millions of Euros, electricity bought in GWh. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01.

Table 6: Effects of the EU ETS on Pollution Control Investment

	(1)	(2)	(3)
	Pollution Control: Air Quality & Climate Change		
	$\Delta \operatorname{arcsinh}$ (Measurement)	$\Delta \operatorname{arcsinh}$ (Integrated)	$\Delta \operatorname{arcsinh}$ (Specific)
Pre-Announcement	0.177 (0.307)	- -	- -
Announcement Period	-0.085 (0.318)	0.041 (0.231)	0.061 (0.233)
Trading Phase I	-0.159 (0.362)	0.975 (0.359)	0.006 (0.334)
Trading Phase II	-0.270 (0.511)	1.079 (0.401)	0.072 (0.341)
Mean in 2001	18.828	56.726	47.008
Observations	17,073	17,507	17,507
Total # of Firms	1,732	2,936	2,936
# of Regulated Firms	128	158	158
Adjusted R ²	0.061	0.037	0.07

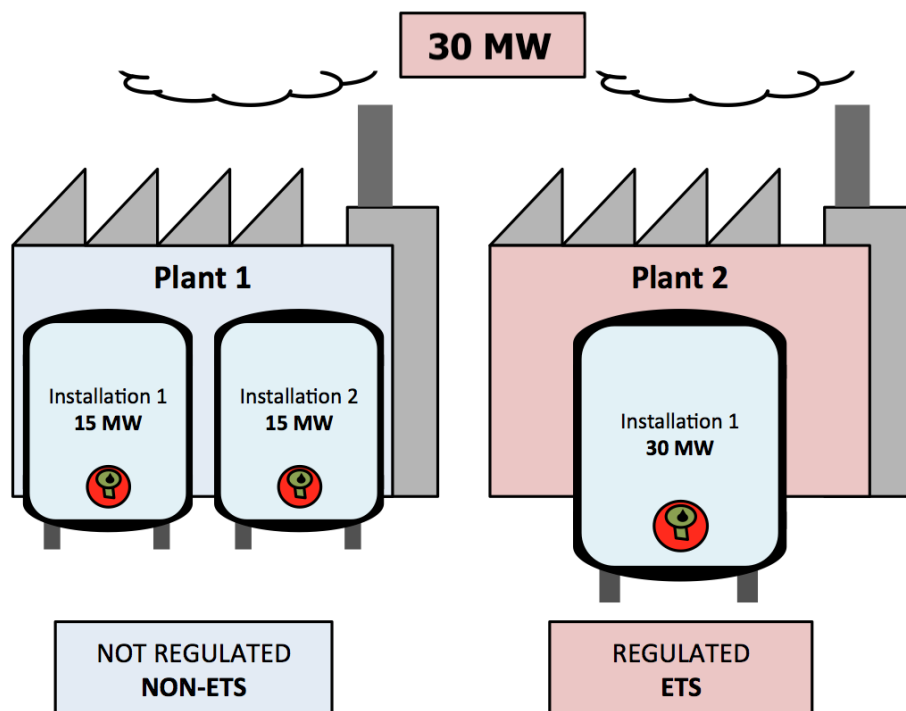
Notes: The dependent variables are: in column 1 { investment into the measurement of emissions; in column 2 { investment in integrated investments made in production processes and machines that are less carbon- or air pollution-intensive than alternatives; in column 3 { investments into specific, 'end-of-pipe' measures to reduce emissions. We take the arc sine of all three variables. These estimates are the result of OLS regressions. They provide the difference between regulated firm and unregulated firm outcomes prior to the announcement of the EU ETS, during the announcement period and during Phase I and Phase II of the EU ETS. Each coefficient represents the difference relative to the year 2001, except in column 1, where it is relative to the year 2000. Means are reported for ETS firms in 2001. Units: Thousands of Euros. Standard errors are clustered at the firm-level. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01.

Table 7: Effects of the EU ETS on Pollution Control Investment (Matched Sample)

	(1)	(2)	(3)
	Pollution Control: Air Quality & Climate Change		
	$\Delta \operatorname{arcsinh}$ (Measurement)	$\Delta \operatorname{arcsinh}$ (Integrated)	$\Delta \operatorname{arcsinh}$ (Specific)
Pre-Announcement	-0.371 (0.530)	- -	- -
Announcement Period	0.154 (0.461)	-0.100 (0.339)	-0.569 (0.500)
Trading Phase I	0.296 (0.634)	0.835 (0.398)	-0.720 (0.709)
Trading Phase II	0.046 (0.821)	0.501 (0.676)	-0.216 (0.503)
Mean in 2001	12.105	25.536	28.924
Observations	1,088	1,000	1,000
# of Regulated Firms	113	142	142
Adjusted R ²	-0.001	0.007	0.010


Notes: The dependent variables are: in column 1 { investment into the measurement of emissions; in column 2 { investment in integrated investments made in production processes and machines that are less carbon- or air pollution-intensive than alternatives; in column 3 { investments into specific, 'end-of-pipe' measures to reduce emissions. We take the arc sine of all three variables. These estimates are the result of OLS regressions, estimated on a matched sample. They provide the difference between regulated firm and unregulated firm outcomes prior to the announcement of the EU ETS, during the announcement period and during Phase I and Phase II of the EU ETS. Each coefficient represents the difference relative to the year 2001, except in column 1, where it is relative to the year 2000. Means are reported for ETS firms in 2001 except in column 1 where it is in 2000. Units: Thousands of Euros. Standard errors are clustered at the firm-level. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01.

Figure 1: Regulation status at the plant level



Notes: The figure depicts an example of how the distribution of a fixed combustion capacity of 30 MW determines the treatment status. Plant 2 operates a single installation with a rated thermal input of 30 MW. Because the participation threshold is at 20 MW, plant 2 must participate in the ETS. Plant 1 has the same total combustion capacity but it operates two installations rated 15 MW each, small enough to not be regulated under the EU ETS.

Figure 2: The Effect of the EU Emissions Trading Scheme on the Environmental and Economic Performance of Firms



(a) CO₂ Emissions

(b) Carbon Intensity




(c) Value Added

(d) Employment

(e) Capital Stock


Notes: The figure shows the effect of the EU ETS on regulated firms, compared to unregulated firms for various outcome variables. All variables are in logs and normalized at the year 2000. Vertical red lines relate to the different phases of the EU ETS. The EU ETS was announced in 2000 and the first phase began in 2005. Phase Two of the EU ETS began in 2008. Standard errors are clustered at the firm-level. In panel a) we observe reductions in emissions in the second phase of the EU ETS. In panel b) we observe a comparable reduction in the emissions intensity of value added. This suggests that the ETS had limited effects on production. Consistent with this we do not observe any negative effects on the economic performance of firms (panels c,d, and e).

Figure 3: The Effect of the EU ETS on Aggregate Emissions Reductions



Sources: Authors calculations based on French microdata and Eurostat data.
 Notes: The black line presents the aggregate time series for industrial emissions in France, measured in millions of tonnes of CO₂. The dark gray line represents counterfactual emissions in the absence of the EU ETS, using our difference-in-differences estimates and assuming that 75% of industrial emissions are regulated. The light gray line represents the level of emissions in 2004 as a benchmark. We estimate that, on average, the EU ETS contributed approximately 26% of the observed emission reductions during this period.

Figure 4: Comparing the EU ETS to other Climate Change Mitigation Policies



Notes: This figure ranks different climate change policies by the estimated cost of reducing a ton of CO₂ in \$2017. The value chosen for the EU ETS is the maximum permit price that was observed during phase II - €29.33 on 1st July 2008. We then convert this to U.S. dollars using the exchange rate on that day and then account for inflation between 2008 and 2017. The maximum cost of reducing a ton of CO₂ was \$52.68. The actual cost was likely far lower, as this is the maximum price at which firms would have been indifferent between reducing emissions and buying permits. Despite this conservative choice, the EU ETS is ranked 7th out of 25. The cost of other policies are taken from [Gillingham & Stock \(2018\)](#). Where multiple estimates exist for the same policy we take the average across all estimates.

Online Appendices – Not for Publication

Contents

A	Descriptive Statistics	i
B	Semi-Parametric Diff-in-Diff Research Design	ii
B.1	Balance and Common Support	ii
B.2	Robustness Tests	vi
B.2.1	Alternative Matching Specifications	vi
B.2.2	Alternative Distance Restrictions	vii
C	Plant-level Results	i
D	Heterogeneity of Treatment Effects with Respect to Permit Allocations	ii
E	Back-of-the-Envelope Calculations	v
E.1	Aggregate Effects of the EU ETS	v
E.2	Extrapolating aggregate emissions impacts to the European Union	vi

A Descriptive Statistics

Table A1: Descriptive Statistics

	(1)	(2)	(3)	(4)	(5)	(6)
	Observations	Mean	St. Dev.	10th perc.	Median	90th perc.
CO ₂	42,733	13.861	88.822	0.165	1.192	14.285
Employment	42,733	450	968	95	241	876
Value Added	42,733	27.563	71.366	3.640	10.350	55.217
Capital Stock	42,733	46.567	13.491	3.342	15.422	96.598
CO ₂ /VA	42,733	0.536	2.283	0.018	0.096	1.173
Total Imports	42,733	22.261	117.621	0.266	4.229	41.537
Carbon Intensive Imports	42,733	0.282	1.874	0.002	0.044	0.484
CO ₂ Gas	42,733	7.048	37.790	0	0.695	10.895
Gas Share	42,733	0.678	0.418	0	0.958	1
Electricity Bought	42,733	21.704	127.091	0.802	4.690	32.665
Electricity Bought/Consumed	42,732	0.991	0.053	0.995	1	1
Pollution Control Investment:						
Measurement	29,748	2.849	28.478	0	0	0.338
Integrated	29,748	12.471	122.366	0	0	1.545
Specific	29,748	11.877	186.062	0	0	0

Notes: Column 1 reports the number of observations (rms x years), Columns 2 and 3 report the mean and standard deviation of each variable. Columns 4 to 6 present the median, 10th percentile and 90th percentile. Units CO₂ and CO₂ from Gas { thousands of tons of CO₂; Value Added { millions of Euros; Employment { full-time equivalent employees; Capital { millions of Euros; CO₂/VA and CO₂ from Gas/VA units { tonnes of CO₂ per thousand Euros of value added; Imports { millions of Euros; Electricity Bought { GWh; Pollution control investment { Thousands of Euros.

B Semi-Parametric Diff-in-Diff Research Design

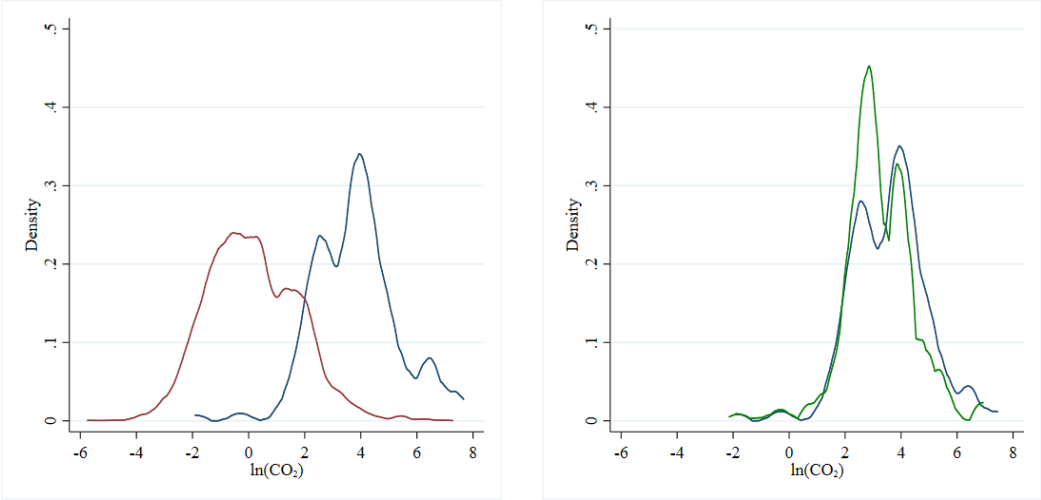
B.1 Balance and Common Support

Table B1: Descriptive Statistics for Regulated and Unregulated Firms

	(1)	(2)	(3)	(4)
	Pre-Match Unregulated	Pre-Match Regulated	Pre-Match Difference	Post-Match Difference
	(Full Sample)	(Full Sample)	(Full Sample)	(Matched Sample)
log (CO ₂)	0.037 (1.684)	3.735 (1.504)	3.698 (0.121)	0.351 (0.174)
log (Employment)	5.466 (0.873)	6.112 (1.312)	0.646 (0.104)	0.281 (0.205)
log (Value Added)	9.260 (1.041)	10.309 (1.423)	1.049 (0.112)	0.215 (0.177)
log (Capital Stock)	9.477 (1.285)	11.212 (1.459)	1.735 (0.116)	0.413 (0.181)
log (CO ₂ /VA)	2.290 (1.554)	4.939 (1.403)	2.649 (0.113)	0.136 (0.195)
log (Total Imports)	14.945 (2.064)	16.251 (2.304)	1.306 (0.183)	0.139 (0.338)
log (Carbon Intensive Imports)	10.015 (2.197)	11.965 (2.218)	1.950 (0.178)	0.380 (0.362)
log (CO ₂ Gas)	7.002 (1.682)	10.310 (1.302)	3.309 (0.111)	0.473 (0.204)
Gas Share	0.221 (0.432)	0.369 (0.364)	0.063 (0.0293)	-0.0762 (0.0542)
log (Electricity Bought)	8.310 (1.371)	10.804 (1.581)	2.494 (0.126)	0.567 (0.224)
Electricity Bought/Consumed	0.992 (0.050)	0.941 (0.154)	-0.0510 (0.0121)	-0.0280 (0.0155)
Observations in year 2000	3,674	163	3,837	298
# of Regulated Firms	0	163	163	149

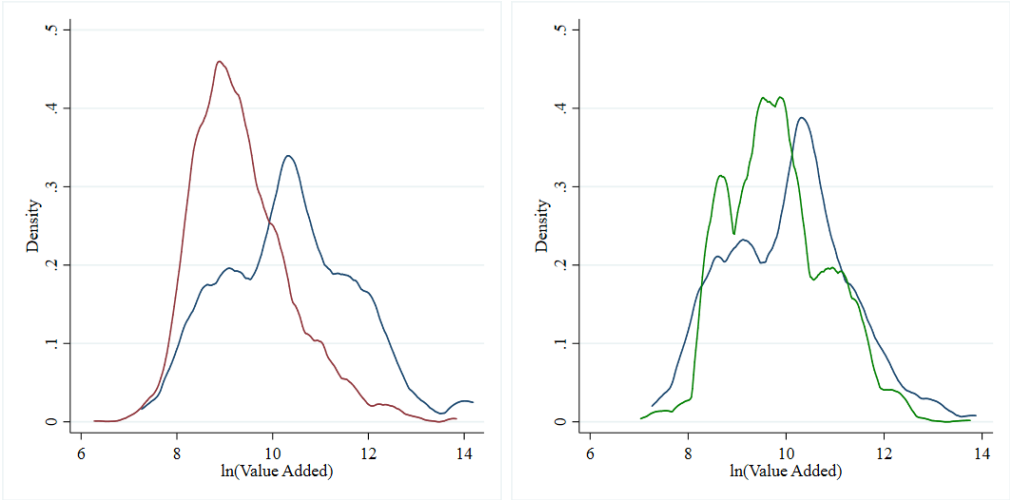
Notes: Columns 1 and 2 report the mean and standard deviation of each variable, respectively for unregulated (control) and regulated (treatment) firms in the year 2000. Reported coefficients in Columns 3 and 4 measure the difference in outcome variables between treatment and control firms in that year. Column 3 presents the average difference between unmatched treatment and control firms. Column 4 presents the average difference between matched treatment and control firms. Standard errors reported in column 3 are robust, and robust and two-way clustered by firm and matching group column 4. We see that the baseline differences are much smaller after matching regulated to unregulated firms. Units (Logarithms of): CO₂ and CO₂ Gas { thousands of tons of CO₂; Value Added { millions of Euros; Employment { full-time equivalent employees; Capital { thousands of Euros; CO₂/VA units { hundred thousands of tonnes of CO₂ per Euros of value added; Imports { Euros; Electricity Bought { MWh. Significance levels are indicated as * 0.10, ** 0.05, *** 0.01.

Figure B1: Density plots showing differences between regulated and unregulated firms in the pre- and post-match samples (CO₂ Emissions, Value Added)



(a) CO₂ (Unmatched)

(b) CO₂ (Matched)




(c) Value Added (Unmatched)


(d) Value Added (Matched)

Notes: The figures report the density plots of log CO₂ emissions and log value added in the year 2000, our base year. In all figures the blue lines represent regulated firms. In the unmatched sample the distribution of each variable for unregulated firms is represented in red. For the matched sample the distribution of each variable for unregulated firms in represented in green.


Figure B2: Density plots showing differences between regulated and unregulated firms in the pre- and post-match samples (Employment and Capital)




(a) Employment (Unmatched)



(b) Employment (Matched)




(c) Capital (Unmatched)



(d) Capital (Matched)

Notes: The figures report the density plots of log employment and log capital in the year 2000, our base year. In all figures the blue lines represent regulated firms. In the unmatched sample the distribution of each variable for unregulated firms is represented in red. For the matched sample the distribution of each variable for unregulated firms in represented in green.

Figure B3: Density plots showing differences between regulated and unregulated firms in the pre- and post-match samples ($\text{CO}_2/\text{Value Added}$)



(c) $\text{CO}_2/\text{Value Added}$ (Unmatched)

(d) $\text{CO}_2/\text{Value Added}$ (Matched)

Notes: The figures report the density plots of log emissions intensity defined as CO_2 emissions/value added in the year 2000, our base year. The blue lines represent regulated firms. In the unmatched sample the distribution for unregulated firms is represented in red. For the matched sample the distribution for unregulated firms is represented in green.

B.2 Robustness Tests

B.2.1 Alternative Matching Specifications

Here we present the regression results exploring the robustness of our emissions reduction results to using different matching variables. Column 1 of Table B2 is identical to Column 1 of Table 2 in the main text. This is our preferred specification and matches firms based on the logarithm of CO₂ in 2000 and the sector they are in. In the other columns of the table we calculate the nearest neighbor match by including additional variables instead of matching only on the logarithm of CO₂ in 2000.

Column 2 matches on the logarithm of CO₂ and the logarithm of the carbon intensity of production (CO₂/VA) in 2000. Column 3 matches on the logarithm of CO₂ and the share of electricity that the firm purchases in 2000. Column 4 matches on the logarithm of CO₂ and the share of emissions that come from non-gas fossil fuels in 2000.

Across all specification, we observe that the estimated reduction in emissions in Phase II is robust to the inclusion of additional matching variables.

Table B2: Matching specifications

	(1)	(2)	(3)	(4)
	$\Delta \log(\text{CO}_2)$	$\Delta \log(\text{CO}_2)$	$\Delta \log(\text{CO}_2)$	$\Delta \log(\text{CO}_2)$
Pre-Announcement	0.016 (0.022)	0.031 (0.024)	0.011 (0.021)	0.022 (0.022)
Announcement Period	0.060 (0.022)	0.021 (0.024)	0.065 (0.023)	0.061 (0.023)
Trading Phase I	-0.002 (0.054)	0.036 (0.057)	-0.002 (0.054)	0.021 (0.055)
Trading Phase II	-0.117 (0.054)	-0.112 (0.057)	-0.115 (0.054)	-0.115 (0.051)
Observations	1,954	2,016	1,954	1,974
Matching Variables	$\ln(\text{CO}_2)$	$\ln(\text{CO}_2)$ & $\ln(\text{CO}_2/\text{VA})$	$\ln(\text{CO}_2)$ & % Elec. Bought/ Consumed	$\ln(\text{CO}_2)$ & % Emissions Gas

Notes: These estimates are the result of OLS regressions, estimated on a matched sample. They provide the difference between regulated firm and unregulated firm outcomes prior to the announcement of the EU ETS, during the announcement period and during Phase I and Phase II of the EU ETS. Each coefficient represents the difference relative to the year 2000. Standard errors are two-way clustered by firm and matching group. Different matching specifications are presented in each column as specified by the line "Matching" where %Elec. Bought/Consumed is the share of electricity consumed that has been bought vs. self-generated. % Emissions Gas is the percentage of GHG emissions due to the burning of gas. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01.

B.2.2 Alternative Distance Restrictions

Table B3 presents an evaluation of how sensitive the baseline matched sample results are to the choice of distance restriction that are imposed when matching treated to control firms. Our main specification from Table 2 is reproduced in column 1. We find that when imposing greater restrictions on the difference between treatment and control firms the significance and magnitude of the reduction in CO₂ emissions in Phase II is reduced.

Table B3: Distance restrictions

	(1)	(2)	(3)	(4)
	$\Delta \log(\text{CO}_2)$	$\Delta \log(\text{CO}_2)$	$\Delta \log(\text{CO}_2)$	$\Delta \log(\text{CO}_2)$
Pre-Announcement	0.016 (0.022)	0.018 (0.022)	0.013 (0.022)	0.010 (0.022)
Announcement Period	0.060 (0.022)	0.061 (0.022)	0.059 (0.022)	0.063 (0.022)
Trading Phase I	-0.002 (0.054)	-0.002 (0.054)	0.000 (0.054)	0.009 (0.054)
Trading Phase II	-0.117 (0.054)	-0.117 (0.054)	-0.108 (0.056)	-0.083 (0.054)
Distance	No Restrictions	99th percentile	95th percentile	90th Percentile
Observations	1,954	1,946	1,896	1,825
# of Regulated Firms	149	148	144	144
Adjusted R ²	0.026	0.027	0.023	0.016

Notes: These estimates are the result of OLS regressions, estimated on a matched sample. They provide the difference between regulated firm and unregulated firm outcomes prior to the announcement of the EU ETS, during the announcement period and during Phase I and Phase II of the EU ETS. Each coefficient represents the difference relative to the year 2000. Distance restrictions between treatment and control firms are imposed at different percentiles. Standard errors are two-way clustered by firm and matching group. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01.

C Plant-level Results

Table C1: The Effect of the EU ETS on Plants Emissions

	(1)	(2)
	$\Delta \log(\text{CO}_2)$	$\Delta \log(\text{CO}_2)$
	Full Sample	Matched Sample
Pre-Announcement	-0.008 (0.024)	-0.006 (0.034)
Announcement Period	-0.018 (0.031)	0.020 (0.026)
Trading Phase I	-0.025 (0.048)	-0.068 (0.050)
Trading Phase II	-0.126 (0.076)	-0.151 (0.075)
Mean in 2000	80.816	79.596
Observations	52,356	6,229
Total Plants	4,944	557
Treated Plants	233	216
Adjusted R ²	0.030	0.052

Notes: These estimates are the result of OLS regressions. The unit of analysis is a plant. They provide the difference between regulated plant and unregulated plant outcomes prior to the announcement of the EU ETS, during the announcement period and during Phase I and Phase II of the EU ETS. Unregulated plants in regulated firms are excluded from the control group. Each coefficient represents the difference relative to the year 2000. Standard errors are clustered by firm. Means are reported for ETS plants in 2000. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01.

D Heterogeneity of Treatment Effects with Respect to Permit Allocations

Our research design recovers the average treatment effects on treated firms. Treatment effects are likely heterogeneous and systematically related to firm attributes that change the intensity of treatment. The firm-specific cost of abating CO₂ emissions is an important attribute, but it is unobservable to us. Differences in treatment intensity may also arise because firms received free permits during the first two phases of the EU ETS. Depending on how generous such permit allocations were, some firms may have found themselves with a surplus and others with a deficit of permits compared to their actual emissions.

In a neo-classical setting without frictions, free permits are equivalent to a lump-sum subsidy in the amount of the value of the total permits received. They should thus not affect marginal production or trading decisions, which are based on the CO₂ price only. In reality, permit allocation could matter because of market power, credit constraints and other frictions, or due to behavioral factors.

To empirically investigate this, one needs a firm-specific measure of permit surplus. We measure permit surplus as the ratio of allocated EUAs to verified emissions in 2005. Figure D1 plots the estimated density function of the permit surplus. We divide firms into two groups – above and below the median surplus – and estimate a variant of equation (1) that includes interactions of the group dummies and the treatment phase dummies. The results are reported in Table D1, where coefficients for both groups represent the full effect for each group to ease interpretation.


The results provide robust evidence that CO₂ abatement among treated firms is driven by those with a below-median surplus (i.e. a rather small surplus or a deficit). The estimates imply an abatement effect of 12.6% in Trading Phase II, which is stronger than the effect obtained in the main results. In contrast, the point estimates in the other group are small and negative yet not statistically significant. Effects on economic performance remain statistically insignificant.³³ However, the positive impact on capital that we found in the main results is also found for the group of abating firms, with a larger and more significant coefficient, in line with our interpretation that abatement occurs via capital upgrades that lower the CO₂ intensity of production. Finally, regressions for the import variables, reported in the last two columns, show no evidence of carbon leakage via imports.

An important caveat is that our measure of heterogeneity – permit surplus – may be endogenous to a firm’s emissions in several ways. For instance, the nominator – permit

³³Differential pre-trends on employment, value-added and CO₂ intensity for below-median firms are statistically significant only at 10%.

allocations – could have been altered as a result of firms that expect to grow lobbying more strongly for a higher number of free permits. Also, the denominator – verified emissions – is in part a result of the firm’s own abatement efforts. We limit the impact of such confounding factors on the results by holding permit surplus fixed at its initial level (i.e. 2005 or, for some firms, the year in which they joined the EU ETS). However, it stands to reason that we cannot rule out selection into groups without a suitable instrumental variable, which we do not have. Therefore, we do not interpret these effects as the causal effect of permits on firm abatement behavior.

Figure D1: Density of Permit Surplus Among Regulated Firms



Notes: The figure depicts the density of the permit surplus, defined as the ratio of allocated permits over verified emissions in the year 2005, or the earliest year available thereafter, for the 163 EU ETS firms of the sample. Data are sourced from the EUTL. Observations of five firms with values above 6 are not depicted on the plot.

Table D1: Permit Allocations and the Heterogenous Effects of the EU ETS (Full Sample)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	log(CO ₂)	log(Value Added)	log(Employment)	log(Capital)	log(CO ₂ /VA)	log(Imports) (Total)	log(Imports) (CO ₂ intensive)
<i>A. Above-median surplus</i>							
Pre-Announcement	-0.031 (0.024)	-0.026 (0.032)	0.012 (0.022)	0.020 (0.021)	-0.005 (0.039)	-0.063 (0.112)	-0.016 (0.134)
Announcement Period	-0.027 (0.030)	-0.026 (0.031)	-0.011 (0.017)	-0.004 (0.019)	-0.002 (0.041)	-0.081 (0.078)	0.017 (0.112)
Trading Phase I	-0.042 (0.043)	-0.095 (0.048)	-0.019 (0.032)	0.011 (0.036)	0.053 (0.049)	-0.137 (0.114)	-0.039 (0.136)
Trading Phase II	-0.013 (0.058)	-0.011 (0.070)	-0.003 (0.050)	0.014 (0.054)	-0.002 (0.066)	-0.099 (0.173)	-0.086 (0.192)
Mean in 2000	86.195	62.707	956	147.148	0.003	75.349	0.749
<i>B. Below-median surplus</i>							
Pre-Announcement	-0.009 (0.028)	0.051 (0.030)	0.037 (0.020)	0.002 (0.024)	-0.060 (0.035)	0.040 (0.106)	0.107 (0.125)
Announcement Period	-0.042 (0.028)	0.028 (0.032)	-0.007 (0.019)	0.001 (0.018)	-0.070 (0.032)	-0.045 (0.087)	-0.186 (0.116)
Trading Phase I	-0.046 (0.041)	0.015 (0.049)	-0.010 (0.034)	0.038 (0.039)	-0.061 (0.049)	-0.079 (0.157)	-0.170 (0.136)
Trading Phase II	-0.134 (0.048)	0.053 (0.060)	0.045 (0.041)	0.102 (0.046)	-0.187 (0.060)	-0.007 (0.154)	-0.201 (0.164)
Mean in 2000	174.249	115.060	1,416	240.527	0.003	142.864	0.991
Observations	42,733	42,733	42,733	42,733	42,733	42,106	41,497
Total # of Firms	3,837	3,837	3,837	3,837	3,837	3,837	3,782
# of Regulated Firms	163	163	163	163	163	163	163
Adjusted R ²	0.030	0.077	0.101	0.450	0.051	0.096	0.154

Notes: These estimates are the result of OLS regressions. They provide the difference between regulated firm and unregulated firm outcomes prior to the announcement of the EU ETS, during the announcement period and during Phase I and Phase II of the EU ETS. Each coefficient represents the difference relative to the year 2000. Coefficients reported in panels A and B are obtained in a single regression for each column, with treatment interactions for whether firms have above or below surplus of permits. These coefficients are displayed as the level effect for a given period plus the interaction term, e.g. Phase II + Phase II (Above or below median Surplus indicator). Means are reported for ETS firms in 2000. Units: CO₂ { thousands of tonnes of CO₂; Value Added { millions of Euros; Capital { millions of Euros; Employment { full-time equivalent employees; CO₂/VA units { thousands of tonnes of CO₂ per thousand Euros of value added; total imports and CO₂ Intensive Imports { millions of Euros. Standard errors are clustered at the firm-level. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01.

E Back-of-the-Envelope Calculations

In the main paper we present several back-of-the-envelope calculations. This section describes the steps that were taken to make these calculations.

E.1 Aggregate Effects of the EU ETS

Our objective is to understand the contribution of the EU ETS to aggregate emission reductions during the period of analysis. Let the variable CO_2 denote aggregate industrial emissions in France, taken from the Eurostat data explorer (<https://ec.europa.eu/eurostat/data/database>). This variable is plotted as a black line in Figure 1.

According to CITEPA,³⁴ 75% of aggregate industrial emissions are regulated. Furthermore, in 2000, 72% of industrial emissions resulted from the combustion of fossil fuels and 28% from industrial processes such as the decarbonation of limestone.³⁵ Since our econometric analysis is based on emissions resulting from combustion only, we refrain from applying the estimated treatment effects to emissions from industrial processes. In order to avoid overestimating the aggregate effects of the EU ETS, we assume that process emissions were not affected by the EU ETS. For lack of more detailed information, we further assume that the 75%-share of ETS regulated emissions applies to combustion and process emissions alike. With this in hand, we can calculate counterfactual emissions in France based on the difference-in-differences estimates from Table 1 as follows:

$$\widehat{CO_2}(t)^{NoETS} = CO_2(t) - 0.75 - 0.72 - CO_2(2000) \left(\sum_{i=1}^4 \mathbf{1}_{ft} 2 \Theta_i g \right)$$

where Θ_i are the treatment phases defined in Appendix 2.2 above. Since only the point estimate for Phase II is statistically significant, we let $\hat{\Theta}_4 = \hat{\Theta}_4$ and $\Theta_1 = \Theta_2 = \Theta_3 = 0$. The variable $\widehat{CO_2}(t)^{NoETS}$ is plotted as a dark gray line in Figure 1.

To approximate the contribution of the EU ETS to aggregate emission reductions we use emissions in the year 2004 as a benchmark. In Figure 1, the level of emissions in 2004 is plotted in light gray.

³⁴The Centre Interprofessionnel Technique d'Étude des Pollutions Atmosphériques - Interprofessional Technical Centre for the Study of Atmospheric Pollutions is a State operator for the French Environment Ministry and a non-profit organisation. Their SECTEN (SECTeurs Économiques et Énergie - Economic and energy sectors) report presents an inventory of GHG emissions and atmospheric pollutants per sector, and is available on <https://www.citepa.org/fr/secten/>. The “Émissions au format Plan Climat” spreadsheet reports, by sector, total and shares of emissions covered or not by the EU ETS.

³⁵CITEPA reports these figures in their UNFCCC inventory available on <https://www.citepa.org/fr/cnucc/>. These shares have been relatively stable over time.

Next, we calculate the difference between observed emissions and emissions in 2004 to get an estimate of the overall change in emissions following the introduction of the EU ETS. The overall reduction in emissions comprises not only the treatment effect of the EU ETS but also the effects of concurrent shocks that may have affected industrial emissions, such as the great recession.

To get a sense of the relative contribution of the EU ETS to total reduction in emissions during this period, we calculate the average difference in emissions, using the following ratio,

$$\frac{\frac{1}{t} \sum_{t=1}^T [\widehat{CO}_2(t)^{NoETS} - CO_2(t)]}{\frac{1}{7} \sum_{t=1}^T [CO_2(2004) - CO_2(t)]}$$

This calculation suggests that approximately 26% of observed emission reductions during the second trading phase can be attributed to the EU ETS, as measured by our difference-in-differences estimates.

E.2 Extrapolating aggregate emissions impacts to the European Union

We repeat the above back-of-envelope calculations using aggregate industrial emissions for the entire EU, obtained from the Eurostat data explorer (<https://ec.europa.eu/eurostat/data/database>). Using the same assumptions and applying our difference-in-differences estimates we calculate that, on average, EU CO₂ emissions would have been 36 million tonnes higher each year in the absence of the EU ETS. Similar to the estimate for France, we calculate that approximately 30% of observed emission reductions since 2004 can be attributed to the EU ETS, as measured by our difference-in-differences estimates.