

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

Jonathan Colmer

University of Virginia, USA

Ralf Martin

International Finance Corporation and Imperial College London, UK

Mirabelle Muûls

Imperial College London, UK; National Bank of Belgium, Belgium

and

Ulrich J. Wagner

University of Mannheim, Germany

First version received August 2021; Editorial decision January 2024; Accepted May 2024 (Eds.)

In theory, market-based regulatory instruments correct market failures at least cost. However, evidence on their efficacy remains scarce. Using administrative data, we estimate that, on average, the European Union Emissions Trading System (EU ETS)—the world’s first and largest market-based climate policy—induced regulated manufacturing firms to reduce carbon dioxide emissions by 14–16% with no detectable contractions in economic activity. We find no evidence of outsourcing to unregulated firms or markets; instead, firms made targeted investments, reducing the emissions intensity of production. These results indicate that the EU ETS induced global emissions reductions, a necessary and sufficient condition for mitigating climate change. We show that the absence of any negative economic effects can be rationalized in a model where pricing the externality induces firms to make fixed-cost investments in energy-saving capital that reduce marginal variable costs.

Key words: Cap-and-trade, Carbon leakage, Investment, Climate policy

JEL codes: Q54, Q58, H23, L50, F18

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 and Oates, 1971; Baumol, 1972; Montgomery, 1972; Tietenberg, 1973; Nordhaus, 1977; Hahn, 1989; Nordhaus, 2001; Burke *et al.*, 2016; Gillingham and 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, then carbon emissions will simply “leak” to unregulated jurisdictions or to unregulated firms or market segments within the same jurisdiction. 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). This is achieved by imposing a cap on the aggregate emissions from more than 12,000 power and manufacturing plants in thirty-one 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 stringency and the extent to which emissions are relocated to unregulated jurisdictions. 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. Figure 1 plots permit prices in the EU ETS during our study period. In trading phase I (from 2005 until 2007), permit prices initially climbed to over €30 but then fell by 50% in April 2006 when evidence came in that the cap was not binding. By the end of 2007, phase I permits were essentially worthless. In contrast, phase II futures prices, which capture the expected stringency of the cap for trading phase II (from 2008 until 2012) remained between 15

1. While there is plenty of disagreement among economists in discussions of policy and government intervention, a preference for market-based regulatory instruments is a point on which economists largely agree. On 17 January 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.




FIGURE 1

EU permit prices during phase I and phase II of the EU ETS

Notes: The figure reports daily average prices of EUA futures (€) between January 2005, the start of phase I, and December 2012, the end of phase II. Reproduced from [Ellerman *et al.* \(2016\)](#) with gracious support by Aleksandar Zaklan.

and 20 euro for 2006 and 2007, before rebounding to € 30 again in 2008. For the remainder of phase II, however, prices declined to between € 8 and € 15. Whether these prices were sufficient to deliver meaningful reductions in regulated emissions, and whether these reductions were offset by increases in unregulated emissions, are empirical questions. We seek to answer these questions using comprehensive administrative data from the French manufacturing sector.

Using a matched difference-in-differences research design, we estimate that the EU ETS induced regulated firms to reduce CO₂ emissions relative to unregulated firms by 14% during trading phase I and by 16% in trading phase II, with no detectable negative effects on economic output or employment. 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 5.4 million tonnes on average between 2005 and 2012, accounting for approximately 28–47% of the aggregate reduction in industrial emissions during this period. We note that our estimates capture the direct effects of the EU ETS on firm behaviour and so likely reflect a lower bound on the aggregate effects of the EU ETS. We do not identify any common firm responses to the EU ETS through market-wide price increases in electricity or other carbon-intensive inputs ([Fabra and Reguant, 2014](#); [Hintermann, 2016](#)).

We also provide evidence indicating that the EU ETS induced global emissions reductions, which is the relevant outcome from the perspective of climate change mitigation. First, as noted, 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 the CO₂ intensity 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 or a change in the composition of emissions. These findings are inconsistent with carbon leakage being a first-order driver of the estimated emissions reductions in this context. Instead, we present evidence that investments in cleaner production processes was the prevailing abatement mechanism among regulated firms.

How could firms reduce emissions without any detectable contraction in economic activity despite the fact that carbon pricing increases input costs? Under standard assumptions, a model of firm production predicts contractions in economic activity alongside reductions in emissions (possibly accompanied by decreasing effects on productivity, cf. [Greenstone *et al.*, 2012](#)). Contrary to this, we find that ETS participation is associated with weakly positive effects on value added, employment, investment, and productivity. One hypothesis is that the ETS induced firms to make investments that increased productivity, offsetting the regulatory costs to the firm, which is also consistent with the slight rebound in emissions, but not emissions intensity, that we observe during the later years of phase II. We present an augmented model of firm production, where firms have the opportunity to switch to an alternative production technology, which requires a fixed switching cost, and also reduces marginal variable costs and weakly increases productivity. In the presence of such a technology, it is no longer clear whether optimal abatement will require the firm to accept higher marginal costs of production or to make a costly, once-and-for-all investment that prevents increasing marginal costs afterwards. Many existing technologies (because they economize on energy) could actually reduce marginal production costs, but their adoption is not always profitable. In our model, firms switch if the present discounted value of doing so exceeds the switching cost. When carbon prices are low, no switching occurs. Compared to a counterfactual without carbon pricing, this case gives rise to the standard prediction that reductions in emissions occur alongside a contraction in firm production, as firms face higher marginal costs. At higher carbon prices, switching occurs, leading to a reduction in emissions, and increase in measured capital. When the “clean” production technology also raises total factor productivity (TFP), then value added, employment, and measured TFP also increase. Our empirical results are most consistent with the case in which firms pay fixed up-front costs to switch into “clean” production technologies that reduce the emissions intensity of production, reduce marginal variable costs, and increase productivity, offsetting the direct costs of carbon pricing.

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 favourably to the marginal abatement costs of many non-market-based regulatory instruments ([Gillingham and 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 behaviour ([Becker and Henderson, 2000](#); [Greenstone, 2002](#); [Fowle *et al.*, 2012](#); [Greenstone *et al.*, 2012](#); [Ryan, 2012](#); [Walker, 2013](#); [Martin *et al.*, 2014a, 2014b](#); [Fowle *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. This is essential to understanding whether the policy was effective at achieving its ultimate objective, which is to reduce global emissions. We also present a new framework for evaluating the economic consequences of environmental regulations on firm behaviour. This framework provides a helpful structure to discipline the interpretation of our empirical results, and provides guidance for future research in this area.

Second, we contribute to a growing empirical literature seeking to understand the effects of the EU ETS itself (see [Martin *et al.*, 2016](#), for a more detailed review). Early studies in this area have been at the country or sector level, which complicates causal inference due to confounding factors ([Ellerman and Feilhaber, 2008](#); [Ellerman *et al.*, 2010](#); [Andersen and Di Maria, 2011](#);

Egenhofer *et al.*, 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 (Jaraite and Di Maria, 2016; Marin *et al.*, 2018; Löschel *et al.*, 2019; Klemetsen *et al.*, 2020; Gerster *et al.*, 2021; Dechezleprêtre *et al.*, 2023). 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 Germany (Gerster *et al.*, 2021) or Lithuania where CO₂ intensity fell (Jaraite and 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.*, 2023). The tightening of the EU ETS in more recent years has come with improvements in the emission efficiency of the biggest emitters (De Jonghe *et al.*, 2020).

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, we not only 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 regulation. In doing so, 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 and Finkelstein, 2020; Hendren and Sprung-Keyser, 2020); however, evaluating the welfare effects of regulations faces a number of theoretical and empirical challenges, given the need to weigh the benefits to society against the costs to firms and workers. Our findings indicate that the EU ETS delivered global emissions reductions with no detectable economic contraction. Understanding the efficacy of government interventions is especially important in the context of mitigating climate change, due to the severity of the problem and due to the limited resources available to tackle it. Through the lens of our model, our findings suggest that the costs associated with decarbonization may only be costly in the transition phase, rather than in the long term. We posit that the emissions reductions induced by the EU ETS likely cost substantially less per tonne of CO₂ than alternative non-market-based regulatory instruments (Gillingham and Stock, 2018).³

In the next section, we describe the design of the EU ETS and our empirical approach. Section 3 describes the data used for analysis. Section 4 presents the main results and Section 5 explores the underlying mechanisms. Section 6 presents back-of-the-envelope calculations that

2. Beyond manufacturing, researchers have estimated the impact of the EU ETS on power plants (Fabra and Reguant, 2014; Zaklan, 2023), on patenting (Calel and Dechezleprêtre, 2016), and on foreign direct investment (Koch and Basse Mama, 2019; Borghesi *et al.*, 2020).

3. This conclusion only holds for the manufacturing sector considered here; a system-wide assessment of abatement costs is beyond the scope of our study. Moreover, as noted by Vogt-Schilb and Hallegatte (2014), command-and-control policies might deliver better results if emissions-reducing investments are subject to strong path dependencies, requiring that expensive abatement investments be made before reaping low-hanging fruit. We thank two anonymous referees for raising these caveats.

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 7 concludes.

2. EVALUATING THE EU ETS

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, since most countries did not explicitly collect this information before it was required for monitoring purposes.⁴ However, pre-implementation data are necessary 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, is explained in detail in Section 3.

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 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 resulted in a significant drop in economic activity, which in turn likely contributed to at least temporary declines in GHG emissions in the EU and around the world. These trends make the evaluation of emissions trading systems 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.

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 direct effects 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 programme for CO₂ emissions.⁵ Polluters regulated under the policy are required to surrender, at the end of each year, one European Union Allowance (EUA) for each tonne of CO₂ equivalent 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 and, since 2008, across trading phases. 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. Allocation of EUAs to polluters is via free allocation or permit

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

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

auctions. During the study period of this paper, 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 behaviour at the margin. Second, we lack a credible strategy to test for a causal effect.

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. While this makes 2005 the first year of actual regulatory treatment, we allow for the possibility that polluters 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-firm level, referred to in the directive as installations. Different criteria are defined for combustion activities on the one hand and other carbon-intensive processes on the other hand.

Participation in the EU ETS is mandatory for combustion installations with a rated thermal input of 20 megawatts (MW) or more. This concerns not only fossil-fuel fired power plants, which are not analysed in this paper, but also industrial plants across a wide range of industries that generate heat, steam, or power on site. Additional industrial installations are included because they specialize in carbon-intensive processes and exceed specific capacity thresholds. Process-based definitions target, among others 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.* emissions from sources that are not owned and not directly controlled by the firm, are not taken into account, nor are electricity imports.

We match French ETS installations listed in the official trading registry to the manufacturing establishments operating them (further detail is presented in Section 3.7). 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 operates at least one ETS plant. We define a time-invariant definition of exposure to the ETS based on whether a firm has ever operated at least one ETS plant during the study period.

The installation-centered, capacity-based participation rules used in the Emissions Trading Directive induce variation in treatment status even among firms of similar size

6. Since 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 therefore, high compliance costs—for decades to come.

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

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

(Calel and Dechezleprêtre, 2016). To see this, consider the case of two firms that operate combustion installations. Both firms have two plants and a total combustion capacity of 30 MW, but the distribution of that capacity across plants gives rise to different treatment assignments. One of Firm 1's plants is treated because it has a rated thermal input of 25 MW, which is above the participation threshold. The other plant has a rated thermal input of 5 MW and is untreated. We define Firm 1 as treated because one of its plants is regulated. Firm 2 is not regulated because it achieves the same total capacity by operating two smaller plants with rated capacity of 15 MW each, 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 plants 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). Nevertheless, we can take advantage of the fact that the participation rules induce variation in treatment status across firms with similar levels of CO₂ emissions 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 set out to identify average treatment effects on the treated at the firm level.

Table 1 presents within-sector differences in pre-treatment characteristics between ETS and non-ETS firms in the year 2000. We see that there are large and significant differences in emissions and production between regulated and unregulated firms. While balance is not required to identify the effects of the ETS using a difference-in-differences estimator, the parallel trends assumption is more likely to hold when baseline differences between the treatment and control group are smaller. The large gaps motivate the creation of a matched analysis sample, which we use in our main analysis. We discuss the matching process below, but note that while some baseline differences remain between treated and control firms they are notably smaller than in the unmatched sample and statistically insignificant in many cases.

2.2. Matched 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. We use a semi-parametric difference-in-differences approach, following Heckman *et al.* (1997, 1998):

$$\begin{aligned} \alpha_{\text{ATT}}^{\text{matched}} &= \mathbb{E}[Y_{it'}(1) - Y_{it'}(0) | X_i, \text{ETS}_i = 1] \\ &= \frac{1}{N_1} \sum_{j \in I_1} \left\{ (Y_{jt'}(1) - Y_{jt'}(0)) - \sum_{k \in I_0} \omega_{jk}(X_j, X_k) \cdot (Y_{kt'}(0) - Y_{kt'}(0)) \right\}, \end{aligned} \quad (2.1)$$

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 ω_{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. Table 1 and Figures A.3 and A.4, Supplementary Material, show that our matching approach, while not perfect, substantially improves the balance and common support between regulated and unregulated firms.

TABLE 1
Descriptive statistics for regulated and unregulated firms

	(1)	(2)	(3)	(4)
	Pre-match unregulated (full sample)	Pre-match regulated (full sample)	Pre-match difference (full sample)	Post-match difference (matched sample)
Log (CO ₂)	−0.043 (1.757)	3.715 (1.527)	3.758*** (0.100)	0.944*** (0.157)
Log (employment)	5.457 (0.873)	6.126 (1.265)	0.668*** (0.0808)	0.135 (0.0993)
Log (value added)	9.242 (1.047)	10.295 (1.361)	1.053*** (0.0872)	0.176 (0.120)
Log (capital stock)	9.449 (1.310)	11.233 (1.534)	1.784*** (0.0987)	0.444*** (0.152)
Log (CO ₂ /VA)	2.228 (1.636)	4.933 (1.395)	2.705*** (0.0915)	0.768*** (0.0936)
Log (total imports)	16.052 (1.401)	17.139 (1.823)	1.087*** (0.117)	−0.0114 (0.222)
Gas share	0.638 (0.440)	0.702 (0.372)	0.0638*** (0.0244)	−0.0647 (0.0592)
Electricity bought share	0.516 (0.247)	0.263 (0.188)	−0.254*** (0.0125)	−0.0375** (0.0171)
Observations in year 2000	3,949	252	4,201	298
# of regulated firms	0	252	252	149

Notes: Columns 1 and 2 report the mean and standard deviation of each variable 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. Robust standard errors reported in column 3. Two-way clustered standard errors (by firm and matching group) are reported in column 4. Units (logarithms of): CO₂—thousands of tonnes of CO₂; value added—thousands of euro; employment—full-time equivalent employees; capital—thousands of euro; CO₂/VA units—hundred thousands of tonnes of CO₂ per euro of value added; imports—euro; gas share—CO₂ from gas/total CO₂; electricity bought share—purchased electricity/total energy consumed in tonnes of oil equivalent. Purchased electricity is converted from MWh to tonnes of oil equivalent using the conversion factor, tonnes of oil equivalent = MWh × 0.086. Significance levels are indicated as **0.05, ***0.01.

In our baseline specification, we implement this approach as a difference-in-differences regression on a matched sample obtained in a one-to-one nearest-neighbour matching. We calculate the difference in average emissions for regulated firms, before and after the introduction of the EU ETS and subtract from this change the difference in average emissions from a matched unregulated firm before and after the introduction of the EU ETS. The regression equation is given by

$$(Y_{j,t} - Y_{j,2000}) - (Y_{k,t} - Y_{k,2000}) = \sum_{\tau=1}^4 \beta_{\tau} \times \mathbb{1}\{t \in \Phi_{\tau}\} + \varepsilon_{j,t}, \quad (2.2)$$

where phases $\{\Phi_{\tau}\}_{\tau=1}^4$ are defined as

- $\Phi_1 = \{1996, \dots, 1999\}$ (pre-announcement period),
- $\Phi_2 = \{2001, \dots, 2004\}$ (announcement period),
- $\Phi_3 = \{2005, \dots, 2007\}$ (trading phase I), and
- $\Phi_4 = \{2008, \dots, 2012\}$ (trading phase II).

The left-hand side of [equation \(2\)](#), Supplementary Material, 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 $\beta_\tau = \alpha_{ATT}^{\text{matched}}$ and provide the effect of the EU ETS on regulated firms in period τ as compared to the matched control firms, and relative to the year 2000.

Matching variables. We match non-ETS firms to ETS firms along a number of dimensions. For each variable we match using data from the year 2000 (the year prior to the announcement of the EU ETS). We match on the CO₂ emissions, value added, employment, capital, emissions intensity, total imports, share of gas in CO₂ emissions, share of consumed energy that comes from purchased electricity, number of plants in the firm, and the 2-digit NCE sector of the firm, which we re-define to reflect the fact that multi-plant firms may engage in multiple activities.⁹ We match exactly on sector to control for sector-specific shocks to the outcome variables that may have occurred after the introduction of the EU ETS. Within a given sector, we use a nearest neighbour using a mahalanobis distance across our matching variables. Our matching variables are chosen to identify a set of comparison firms that are similar in terms of their environmental characteristics (emissions, emissions intensity), their production function (value added, labour, and capital), the composition of emissions and energy use (gas share and electricity share), and their exposure to trade (imports). We do not match on pre-treatment trends in our baseline specification. Instead, we let the data speak to the validity of the assumption that pre-treatment trends in the outcome variables are parallel. Column 4 of [Table 1](#) shows that the post-match difference in baseline characteristics is substantially smaller than the pre-match difference (column 3). While remaining statistically significant, the gap in emissions, capital, and emissions intensity is 75% smaller than the pre-match difference. The gap in the share of energy consumed that comes from purchased electricity is 85% smaller. There is no statistically significant or economically meaningful post-match difference in value added, employment, the composition of emissions, or imports.

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 and Spiess \(2022\)](#) show 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.

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 fulfils the matching criteria. Given the bias–variance trade-off, we give priority to minimizing bias and use matching with replacement in our main specification. Drawing inspiration from [Abadie and Spiess \(2022\)](#), we use a two-way cluster adjustment to try to address bias in the estimation of the variance. The first

9. We define a new sector variable SUPERNCE 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 NCE 12, then the SUPERNCE is 12 and the firm would be matched to a control firm in the same sector (with SUPERNCE 12). In contrast, for a firm with one plant producing in NCE 12 and another one in NCE 17, we define SUPERNCE to be 1217 and match it to a control firm within SUPERNCE 1217 (where the ordering of sectoral codes does not matter, *e.g.* SUPERNCE “1217” is equivalent to SUPERNCE “1712”).

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 statistical inference. Our adjustment collapses to the solution presented in [Abadie and Spiess \(2022\)](#) when each treatment firm is matched to a unique control firm. In this case, the second cluster becomes redundant. In Appendix B.2, Supplementary Material, we show that our results are robust to using matching without replacement; however, the sample size is smaller and balance between treatment and control firms is worse, consistent with the bias–variance trade-off. As such, we prefer matching with replacement in our baseline specification. The similarity in statistical inference between matching without replacement following [Abadie and Spiess \(2022\)](#) and matching with replacement, applying the two-way cluster adjustment, is encouraging. Inferences are unchanged if we two-way block bootstrap our standard errors. Standard errors are notably smaller when we match with replacement and only cluster standard errors at the firm level, consistent with the insight from [Abadie and Spiess \(2022\)](#), which indicates that there is value added to the two-way adjustment.

2.3. Identification assumptions

Our econometric approach assumes 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 this parallel trends assumption is plausible when evaluating the effects of the EU ETS using pairs of similar firms matched within narrowly defined sectors. To make this argument, it is helpful to distinguish between two potential violations of the parallel trends assumption. First, treated and control firms could be on different trajectories already before the launch of the ETS. Second, other contemporaneous shocks may differentially affect the trajectories of treated and control firms. Either violation would lead to biased inferences about the effect of the ETS. While neither assumption is testable, analysis can help to evaluate whether the violations are likely to be a first-order concern. For example, for observable characteristics, we should not see any differential trends between regulated and unregulated firms prior to the introduction of the ETS. Concerns related to other shocks depend on whether they coincided with the EU ETS and whether treated and control firms were affected differently. Where possible, we engage in additional analyses to help increase the credibility of our research design.

Potential violations may arise from overlapping energy policies and economic fluctuations that occurred during the treatment period. The former include energy taxes, subsidies for renewable energy, and energy efficiency targets. The latter prominently features the Great Recession, which began in the first year of the second trading phase. We engage seriously with the concern that these policies and events may have affected regulated firms differently to matched control firms. Relevant energy policies are reviewed in Appendix B.4, Supplementary Material, with a focus on whether they have different implications for ETS and non-ETS firms, after matching. For policies that pre-dated the ETS, we would expect divergent pre-trends if the policies had any differential effect on regulated firms. Any confounding effect of subsidies for renewables should lead to a differential effect on electricity generation. We do not find any evidence of this.

The Great Recession might confound estimated treatment effects in phase II of the ETS if the economic downturn or the subsequent recovery had a differential effect on firms that have characteristics associated with ETS participation. For example, the size differences in the unmatched sample, highlighted in Table 1, could lead us to underestimate/overstate emissions reductions if untreated small firms were more/less affected during the Great Recession than the larger, capital-intensive, firms that are treated. Matching on a broad set of covariates helps to

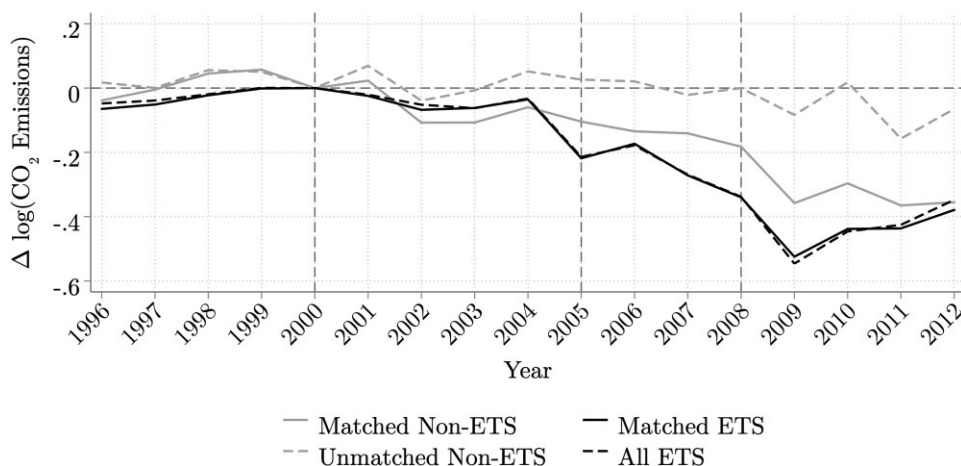


FIGURE 2

Trends in CO₂ emissions by group of firms

Notes: The figure reports average trends in (log) CO₂ emissions relative to the base year 2000 for various groupings of firms in our dataset: all ETS firms, ETS firms for which we can find a non-ETS control firm (matched ETS firms), those control group firms (matched non-ETS), and firms that are not in the ETS nor the matched control group (unmatched non-ETS).

reduce the potential for this issue by minimizing differences in firm size, access to capital, scale economies, and other potentially relevant differences, to the degree that they are captured by the observable dimensions that we match on. To further explore the potential contribution of the Great Recession, we construct geographic and industry-level measures of exposure and explore the robustness of our findings to accounting for these measures. Appendix B.3, Supplementary Material, discusses our estimation strategies and shows that our results in Section 4 are robust to accounting for differential exposure to these measures of the Great Recession.

Further descriptive support for this conclusion is presented in Figure 2, which plots raw trends in CO₂ emissions, by treatment status, for matched and unmatched firms. While there is a clear fall in emissions following the Great Recession in 2008, this drop appears to occur in a near-parallel way for treated firms and matched control firms. In 2011 and 2012, we see more of an uptick in emissions for regulated firms in the raw data, which would lead us to underestimate the effect of the EU ETS if regulated firms were differentially affected during the recovery of the Great Recession.

Figure 2 also provides more general support for the parallel trends assumption prior to the introduction of the EU ETS in 2004. We see that the trajectory of emissions for matched non-ETS firms follows ETS firms closely until 2005, when permit trading begins. At this point, the emissions of regulated firms drop sharply and remain lower throughout the post-treatment period. The trajectory of emissions for unmatched non-ETS firms follows less closely prior to the introduction of the policy, although even in the raw data the differences are not substantial, with deviations concentrated in the announcement period between 2001 and 2004. The closer mapping between matched non-ETS firms and ETS firms provides further support for the use of matched control firms as a counterfactual for treated firms.

In addition to the parallel trends assumption, we must also assume that there are no spillovers between regulated and unregulated firms. We internalize within-firm spillovers by estimating the effects of the EU ETS at the firm level. We cannot, however, rule out the potential for spillovers

between firms. Such spillovers may take the form of emissions leaking from regulated to unregulated firms. We directly evaluate the potential for spillovers as part of our analysis, and find little evidence to suggest that they are of first-order concern in this context.

3. DATA

This section details the different data 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 trading phases I and II. The data are obtained from various sources.¹⁰

3.1. Energy and emissions data

We obtain detailed fuel use data from the Annual Survey of Industrial Energy Consumption (EACEI), a survey conducted annually by the French National Institute of Statistics and Economic Studies (INSEE).¹¹

The survey provides quantities and values of energy consumed by fuel type—broadly speaking, electricity, steam, fossil fuels, and biofuels.¹² Other variables available in the survey include the geographical location of each establishment and their sectoral NCE 2-digit classification. The NCE is the designated French statistical nomenclature of activity for the study of energy production and consumption.¹³

Having reliable data on CO₂ emissions is of central importance to our study. We calculate 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).¹⁴ Consequently, a firm will only be in our core dataset if it reports detailed energy consumption data under the EACEI, as detailed further in Section 3.7 and Appendix A.1, Supplementary Material. The sampling frame for the EACEI includes all French manufacturing establishments.¹⁵ The response rate is close to 90%. This speaks to the high representativeness of the data, 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.

10. Firm- 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.

11. EACEI is the French acronym for *Enquête annuelle sur les consommations d'énergie dans l'industrie*. INSEE stands for *Institut National de la Statistique et des Études Économiques*. Until 2007, the survey was carried out by the statistical service of the Ministry of Industry, SESSI—*Service d'Études et Statistiques de l'Industrie*.

12. Information for the following fuel types is requested from the surveyed firms: electricity (bought, auto-produced—from thermal or non-thermal process—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 by-product of the chemical decomposition of wood for making paper pulp), wood and its by-products, special renewable fuels, special non-renewable fuels.

13. The NCE is the French acronym for the *Nomenclature d'activités économiques pour l'étude des livraisons et Consommations d'Énergie* and can be put in correspondence with the French NACE rev.2-equivalent NAF classification. https://www.insee.fr/fr/statistiques/fichier/3364874/irecoeacei16_correspondance_NCE_NAF-1.pdf. NACE is the Statistical Classification of Economic Activities in the European Community (for the French term “Nomenclature statistique des Activités économiques dans la Communauté Européenne”). NAF stands for *Nomenclature d'activités française*.

14. ADEME is the French acronym for *Agence de l'Environnement et de la Maîtrise de l'Énergie*. EU ETS participants in France are required to use the ADEME's conversion factors when reporting their emissions.

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

Slightly different sampling weights were applied before and after 2007, but the industrial coverage remained constant, including all manufacturing except the sectors of energy production, agri-food, and sawmills. Around 12,000 establishments are drawn for the sample each year and it includes (i) all industrial establishments with twenty employees or more in the most energy consuming sectors¹⁶; (ii) all establishments with more than ten employees in the manufacturing of industrial gases sector; (iii) all establishments with more than 250 employees on the 31st of December of that year; and (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 manufacturing firms, it draws on a database designed to provide a representative sample, 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 and the resulting dataset we exploit is drawn from the database covering the years 1994–2007.¹⁷ For the years 2008–12, the successor system is called ESANE with the resulting dataset FARE.¹⁸ 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 labour 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 1995–2012 are obtained from French Customs (DGDDI).²⁰ 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, imports 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) (Bergounhon *et al.*, 2018). To harmonize across different thresholds, we set import levels to the highest threshold in the ETS years; *i.e.* €2.3 millions. Since all ETS 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.

16. Manufacture of bricks, tiles, and construction products, in baked clay; manufacture of cement; Manufacture of lime and plaster.

17. SUSE is the French acronym for *Système Unifié de Statistique d'Entreprises*. FICUS stands for *Fichier Complet Unifié de SUSE*.

18. ESANE stands for *Elaboration des Statistiques Annuelles d'Entreprises* and FARE stands for *Fichier Approché des Résultats d'ESANE*.

19. Only observations with non-missing values for employment, value added, emissions, and capital are retained.

20. DGDDI stands for *Direction Générale des Douanes et Droits Indirects*.

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.²¹ Following 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 Nomenclatures²² to obtain the value of imports of goods from a given sector. Multiplying these values with the sector's carbon intensity and aggregating across sectors 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).²³ 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 aimed at reducing air pollution, broadly defined. They are split between (a) investments made to “measure” air and GHG pollution, (b) “integrated” investments made in production processes and machines 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 euro. In estimating (b), the “integrated” investment, respondents are 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) are available from 1996 onwards. 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, $\text{arcsinh}y_{it} = \ln(y_{it} + \sqrt{y_{it}^2 + 1})$. This 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.

21. Cf. 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.

22. This can be accessed at: <https://ec.europa.eu/eurostat/ramon>.

23. In French: *Enquête sur les investissements et les dépenses courantes pour protéger l'environnement*. See Appendix C.1, Supplementary Material, for more information.

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.²⁴ A pollution right in the EU ETS is called a 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.

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 since all four datasets use unique identifiers for firms (SIREN) and plants (SIRET).²⁵ As described in Appendix A, Supplementary Material, linking the EACEI to FICUS/FARE and trade data leads to a sample of 4,201 firms emitting a total of 61.4 million tonnes of CO₂ in 2000, which represents 79.3% of aggregate industrial emissions from combustion of fossil fuels 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 EUTL is managed by Caisse des Dépôts. The latter institution provides a link between the permit identifier (Gestion Informatique des Données des Installations Classées [GIDIC] code) from the national registry and the SIREN identifier from INSEE, allowing for the linking of the EUTL data to the business data. Out of the 4,201 firms, 252 are part of the EU ETS. The main variables are summarized in Appendix Table B.1, Supplementary Material. Appendix Figure A.1, Supplementary Material, provides a visual summary of all the steps involved in the construction of the final sample from the raw data. Comparing emissions computed on the basis of the EACEI to those reported in the EUTL confirms their consistency. Appendix A.1, Supplementary Material, illustrates the 0.96 correlation between these measures. We graphically represent this relationship using a Quantile-Quantile plot (Figure A.2, Supplementary Material).

We reiterate that the policy is not randomly assigned across firms. On average, ETS firms are larger than non-ETS firms in terms of employment, value added, capital, and imports (cf. Table 1). ETS firms also emit more CO₂ emissions and are more carbon intensive. These differences motivate the matching approach discussed in Section 2.1, which substantially reduces baseline differences.

24. When 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, 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.

25. SIREN is the French acronym for *Système d'Identification du Répertoire des Entreprises*. To be precise, plants in the EACEI and Antipol are 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.

TABLE 2
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	0.028 (0.021)	0.009 (0.039)	0.002 (0.025)	-0.012 (0.025)	0.022 (0.037)
Announcement period	0.014 (0.025)	0.014 (0.040)	-0.002 (0.019)	0.014 (0.021)	0.013 (0.034)
Trading phase I	-0.140** (0.057)	-0.050 (0.085)	-0.002 (0.036)	0.083* (0.046)	-0.099 (0.068)
Trading phase II	-0.163** (0.075)	0.097 (0.079)	0.046 (0.050)	0.105* (0.060)	-0.174** (0.075)
Mean in 2000	82.107	55.600	684.215	132.919	3.398
Observations	2,348	2,348	2,348	2,348	2,348

Notes: This table presents estimates from OLS regressions, estimated on a matched sample. Standard errors are clustered in two ways, at the firm level and at the matching group level. Each estimate reflects the difference between regulated firm and unregulated firm outcomes relative to the year 2000. We present estimates for four time periods: prior to the announcement of the EU ETS, during the announcement period and during phase I and phase II of the EU ETS. Means are reported for ETS firms in 2000. Units: CO₂—thousands of tonnes of CO₂; value added—millions of euro; employment—full-time equivalent employees; capital—millions of euro; CO₂/VA units—thousands of tonnes of CO₂ per euro of value added. Significance levels are indicated as *0.10, **0.05.

4. RESULTS

4.1. *Main outcomes*

Table 2 presents our main results. We estimate that, on average, regulated firms reduced emissions by 14% ($p < 0.05$) during trading phase I and by 16.3% ($p < 0.05$) during trading phase II. We fail to reject the null hypothesis that the EU ETS had no effect on the economic performance of firms, as measured by value added or the number of employees. With lower confidence than the emissions results, we estimate that regulated firms increased capital investments during trading phase I (8.3%, $p < 0.1$) and trading phase II (10.5%, $p < 0.1$). Finally, we estimate, consistent with the absence of any economic contraction, that regulated firms reduced the emissions intensity of value added during trading phase II (-17.4%, $p < 0.01$). We estimate a 10% reduction in the emissions intensity of output in phase I but it is not statistically significant at conventional levels. We do not estimate any differential effects between the announcement and implementation of the EU ETS.²⁶

As previously discussed, a key assumption that is 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—the parallel trends assumption. The raw data presented in Figure 2 provides initial support for this assumption. In further support of the parallel trends assumption, we do not estimate any statistically or economically meaningful differences between regulated and unregulated firms prior to the announcement or implementation of the EU ETS. Figure 3 presents a visual representation of these findings. However, we know that there are limitations to evaluating parallel trends based on pre-treatment differences (Freyaldenhoven *et al.*, 2019; Roth, 2022; Rambachan and Roth, 2023). Following Rambachan and Roth (2023), we engage in sensitivity analysis. Instead of imposing that the parallel trends assumption holds exactly, we

26. For the remainder of our results, we present average pre-ETS effects in our results tables. We continue to separately present pre-announcement period and announcement period estimates in robustness tests and sensitivity analysis.

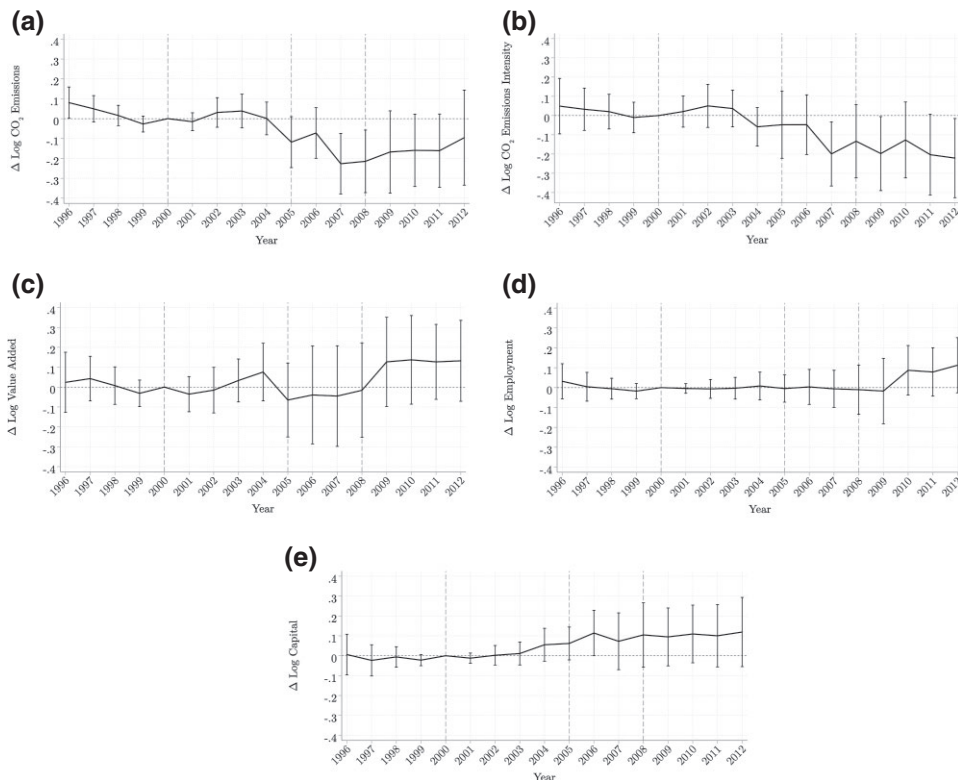


FIGURE 3

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, and (e) capital stock

Notes: These figures present estimates from OLS regressions, estimated on a matched sample. Standard errors are clustered in two ways, at the firm level and at the matching group level. All variables are in logs and normalized at the year 2000. Vertical dashed 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 II of the EU ETS began in 2008. Standard errors are two-way clustered at the firm and match group level.

bound how large post-treatment violations of parallel trends could be before inference “breaks down”. This is formalized by imposing that the post-treatment violation of parallel trends be no more than a constant, \bar{M} , larger than the maximum violation of parallel trends in the pre-treatment period. A value of $\bar{M} = 1$, for example, imposes that the post-treatment violation of parallel trends be no larger than the worst pre-treatment violation of parallel trends (accounting for statistical and identification uncertainty in our event-study estimates).²⁷ For the estimated reduction in emissions, estimated breakdown values are $\bar{M} = 1.7$ for 2007 and 1.3 for 2008.²⁸ Consequently, our conclusion of a significant reduction in emissions depends on whether we are willing to assume that post-treatment violations of parallel trends be no more than 1.3 times as large as the maximal pre-treatment violation. Based on our pre-treatment estimates, differential

27. An alternative framing is to say that post-treatment violations of parallel trends cannot deviate “too much” from a linear extrapolation of the pre-trend; *i.e.* the slope of the pre-trend cannot change by more than “M” across consecutive periods. Imposing a smoothness restriction of $M = 0$ would imply that the counterfactual difference in trends is exactly linear. Larger values of M , by contrast, allow for more non-linear deviations from the pre-trend.

28. Estimates for the other post-treatment years are not statistically significant at the 5% level and so the original 95% confidence already includes zero.

reductions in emissions from other shocks can account for up to 46% of the estimated effect before our inference starts to “break down”. We explore the potential for such violations in the following section. While we cannot rule out violations of parallel trends, these sensitivity tests make clear what must be assumed in order to draw causal inferences. Figure 3 also provides an opportunity to explore dynamics. We estimate an immediate reduction in emissions following the implementation of the ETS in 2005, with the largest reduction in emissions occurring towards the end of phase I and the start of phase II. We estimate a slight reversal of emissions in 2006, which may have arisen due to increased uncertainty about the future stringency of the ETS when it was discovered in April 2006 that the cap was no longer binding for phase I; *i.e.* firms had sufficient permits to remain compliant. This news initially depressed phase II futures prices (Figure 1) and, speculatively, could have delayed some investments until 2007 when prices rebounded. While emissions remained meaningfully below pre-implementation levels throughout phase II, the reductions appear to attenuate over time. In our discussion of mechanisms below, we present a model to reconcile our full set of results. In the model, firms have the option of staying with their current technology and paying higher marginal variable costs, resulting in a contraction, or paying an up-front fixed cost in emissions-saving investments that in turn reduces marginal variable costs. Consistent with technology switching, we observe an initial contraction in phase I, coinciding with the increase in capital investments, followed by a relative expansion in economic activity. While this relative expansion may have attenuated emissions reductions towards the end of phase II, we estimate persistent reduction in the emissions intensity of value added across both phases.

4.2. *Robustness tests*

Our main results are robust to a broad range of alternative specifications and robustness tests. We summarize those findings here and refer the reader to Appendix B, Supplementary Material, for the full set of results.

Table B.4, Supplementary Material, shows that our results are robust to including fewer matching controls (column 2), to matching without replacement (column 3), to increasing the number of nearest neighbours (columns 4–7), to imposing common support on emissions (column 8), and to matching on covariates for all pre-treatment years between 1996 and 2000 (column 9). Table B.5, Supplementary Material, presents post-match baseline differences between treatment and control firms for each matching specification. Compared to our baseline specification, only imposing common support on emissions (column 8) results in smaller differences. In Table B.6, Supplementary Material, we impose increasingly stringent calliper restrictions on the matching distance between treatment and control firms. Our results remain statistically significant until we drop more than 10% of the treated firms with the largest difference; however, even when we drop 25% of treated firms, the results remain qualitatively robust. In column (6), we use the cardinality matching algorithm introduced by Zubizarreta *et al.* (2014). We present it alongside the calliper restrictions because it restricts the data to a sample of matched treatment firms where the differences in matched covariates are no larger than 0.05. As can be seen in Table B.7, Supplementary Material, this results in balance between treatment and control firms, but delivers a sample of firms that is almost half the size of our baseline specification sample. We do not estimate any statistically significant effects, although the estimated emissions reductions in phase I remain qualitatively robust.

More substantive concerns relate to the potential for overlapping policies and events that could differentially affect regulated firms, confounding our interpretation. The absence of any effects on economic performance helps to mitigate the concern that the estimated effects might be confounded by differential reactions to the Great Recession between treated and control

groups. If this were the case, we would expect to see differential responses in economic outcomes as well as environmental outcomes. It is possible that the Great Recession had a differential negative effect on non-ETS firms, offsetting any negative effects of the ETS on regulated firms. In this case, our estimated reductions in emissions would represent a lower bound on the effect of the ETS during the phase II trading period.

A more direct way of assessing confounding effects of the Great Recession is by directly controlling for its effects in the regression. In Appendix B.3, Supplementary Material, we show that this is possible in a straightforward modification to the estimation framework. The intuition behind our approach is that if the recession shocks of treated and matched controls are observed and included in the regression, they can no longer confound the estimated ETS coefficient. While firm-specific recession shocks are unobserved, we can proxy for them using suitable spatial and sectoral measures of unemployment changes between 2008 and 2009. We include these variables separately for treated and control firms when re-estimating our main results. Table B.8, Supplementary Material, shows that the inclusion of these variables has no effect on our phase I treatment effect as should be expected. In phase II, the coefficient is slightly attenuated from -16.3% in our main results to -14.5% . These results lend further credibility to our identifying assumption that the Great Recession did not have differential effects on ETS firms.

A second concern is that other policies may confound the interpretation of our estimates. The EU ETS was not implemented in isolation but in a policy context marked by a commitment by the EU to reduce emissions, with the signing of the Kyoto Protocol in 1997. Under the EU Burden-Sharing Agreement, France was called upon to implement policies in addition to the EU ETS to contribute its fair share to the EU-wide abatement target. Such overlapping policies included energy taxes, subsidies for renewable energy, and the promotion of energy efficiency.

Appendix B.4, Supplementary Material, provides more detail on these policies and explains how differences in the timing of when policies were introduced compared to the EU ETS can be exploited to draw inferences about their empirical contribution to our results. For example, we show that feed-in tariffs for electricity from renewable and small co-generation plants did not affect firms differentially. We conclude that overlapping energy and climate policies in France were unlikely to drive the sizable and robust emissions reductions we estimate in Table 2.

Beyond the Great Recession and introduction of other energy policies, our study spans a time when France is going through a broader process of de-industrialization. Firm exit may have contributed to secular declines in CO₂ emissions. Due to data limitations, we are unable to directly evaluate firm exit. We therefore abstract from firm exit and analyse a balanced sample of firms observed in each one of the four periods. To the degree that the EU ETS induced firms to exit our sample before phase II, our estimated emissions reductions represent a lower bound of the total effect of the EU ETS on industrial emissions. In Appendix B.5, Supplementary Material, however, we provide evidence that there is no differential attrition by ETS firms when constructing our balanced sample.

5. MECHANISMS

Our findings indicate that the EU ETS induced regulated firms to reduce emissions with no detectable effects on economic performance, leading to a reduction in the emissions intensity of production. In this section, we investigate the mechanisms that drive these results.

5.1. *Leakage*

While we estimate that the ETS is associated with reductions in the emissions of regulated firms, what matters for climate change mitigation is whether the ETS reduced global emissions.

Regulated firms may have cut emissions by outsourcing carbon-intensive elements of the value chain to unregulated firms or markets. Carbon leakage threatens the efficacy of the ETS by limiting, or even reversing, the effect on global emissions. For example, if the emissions intensity of production in upstream facilities is higher than in regulated facilities, global emissions could increase as a consequence of the policy.

To assess the efficacy of the EU ETS as a climate policy instrument, it is therefore important to understand 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, *i.e.* by outsourcing 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 such reduction. Moreover, regression results reported in columns 1 and 2 of Table 3 show that there is no statistically significant association between the EU ETS and the importing behaviour of regulated firms; however, the coefficient on total imports in phase II would imply a 4.5% increase in imports if taken at face value.

We bound the potential contribution of imports to our reduction in emissions by combining a naive estimate of the elasticity between emissions intensity and imports, -0.097 , with an upper bound of the increase in total imports (18.6%).²⁹ We calculate that increased imports in phase II could account for at most a 1.8% reduction in emissions intensity, accounting for at most 10% of the estimate in trading phase II. Our collective findings on value added and imports, alongside back-of-the-envelope calculations, provide little evidence to indicate that is likely to be a major driver of our estimated emissions reductions.

The second potential 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 also see negative effects of the EU ETS on at least one economic variable, such as value added, employment, or investment. Instead, however, we estimate insignificant effects on employment and value added, and positive effects on capital investment. Apart from mitigating concerns about leakage, this result is useful as 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 that reallocates market share from regulated to unregulated firms. This would violate the stable unit treatment value assumption (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. The same effect should also 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

29. The elasticity between emissions intensity and imports is estimated using a bivariate OLS regression of the form $\log(\text{CO}_2/\text{value added}_{it}) = \alpha + \beta \log(\text{total imports}_{it}) + \epsilon_{it}$. We estimate the elasticity using all firms in years prior to the EU ETS. The inclusion of firm and sector-year fixed effects attenuates the estimate to -0.022 . The upper bound estimate for the increase in total imports is calculated as $4.5\% + 1.96 \times 7.2\%$.

TABLE 3
Exploring mechanisms

	Import responses		Fuel-mix responses		Pollution control investments			Productivity
	(1) $\Delta \log(\text{imports})$ total	(2) $\Delta \log(\text{imports})$ CO_2 intensive	(3) Δ gas share in CO_2	(4) Δ electricity bought share	(5) Δasinh^{-1} (measurement)	(6) Δasinh^{-1} (integrated)	(7) Δasinh^{-1} (specific)	(8) $\Delta \log(\text{TFRP})$
Pre-ETS	-0.046 (0.033)	-0.050 (0.062)	0.000 (0.004)	0.000 (0.002)	0.230 (0.222)	-0.024 (0.152)	0.108 (0.149)	-0.011 (0.019)
Trading phase I	0.019 (0.060)	0.027 (0.118)	-0.009 (0.017)	-0.004 (0.006)	0.446 (0.327)	0.728*** (0.228)	0.356* (0.202)	-0.026 (0.051)
Trading phase II	0.045 (0.072)	-0.026 (0.134)	-0.018 (0.029)	-0.003 (0.007)	-0.013 (0.310)	0.684** (0.277)	0.404*** (0.146)	0.049 (0.054)
Mean in 2000	156.847	0.752	0.708	0.269	11.51	12.775	15.296	8.180
Observations	2,298	2,267	2,348	2,348	1,237	1,419	1,419	2,270

Notes: These estimates are the result of OLS regressions, estimated on a matched sample. Standard errors are clustered two ways: at the firm level and the matching group level. Each estimate reflects the difference between regulated firm and unregulated firm outcomes prior to implementation of the ETS and during phase I and phase II of the EU ETS. Each coefficient represents the difference relative to the year 2000. Means are reported for ETS firms in 2000 (2001 for columns 6 and 7). Units: total imports and CO_2 intensive imports are measured in millions of euro; gas share of emissions and electricity share of total energy consumed are shares between 0 and 1; investment into the measurement of emissions is measured in thousands of euro; investment in integrated investments made in production processes and machines that are less carbon- or air pollution-intensive than alternative is measured in thousands of euro; investment into specific, "end-of-pipe" measures to reduce emissions is measured in thousands of euro. Significance levels are indicated as *0.10, **0.05, ***0.01.

explain estimated emissions reductions at the firm level. Our estimates are net of any within-firm leakage.³⁰

5.2. *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 from investments in technology that allows fuel switching). Our data allow us to explore these different channels of abatement.

In column 3 of Table 3, we estimate that there was no change in the share of natural gas in total CO₂ emissions. 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. In column 4, we estimate no significant change in the share of electricity procured in firms' total energy use. Firms 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. In Table B.9, Supplementary Material, we estimate no extensive-margin adjustments to on-site generation from conventional and renewable sources. 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.³¹

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. Columns 5–7 in Table 3 provide further evidence in support of this hypothesis using data on pollution control investments for a sub-sample of firms in our sample. 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 6). In column 5, we do not estimate any differential impacts on investments into the measurement of emissions, which in any case is not needed for CO₂, given the ease of input-based accounting. We estimate smaller impacts on investments into specific, “end-of-pipe” measures to reduce emissions, which are not yet easily available for CO₂ at a commercial scale. A caveat to this analysis is that 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. We do not estimate any differential effect prior to the introduction of the ETS for these variables, and for measurement investments we do not estimate any differential effect in the pre-announcement phase.

30. Of all regulated firms, 40% have unregulated CO₂ emissions. In Table A.2, Supplementary Material, we document that the share of total emissions that are regulated is very high in all sectors. In the Pharmaceuticals sector, which has the lowest average share of total emissions that are regulated, 68.22% of emissions are regulated. On average, 88% of emissions in regulated firms are covered by the ETS.

31. It 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, and the remaining 21%—including the marginal generator—is likely to have been produced by power plants under the EU ETS cap.

A review of the metadata of the Antipol survey (see [Section C.1](#), Supplementary Material) provides additional details about the survey but does not provide much detail about the types of investments that firms make. To gain further insight, we take advantage of as-yet unused data from interviews conducted in 2009 with the managers of 140 French manufacturing firms, 92 of which participate in the EU ETS (see [Martin *et al.*, 2014b](#), for details about the data collection). In [Appendix C.2](#), Supplementary Material, we explore responses to interview questions pertaining to measures that were implemented at the production site to reduce CO₂ emissions. Managers were asked “Can you tell me what measures you have adopted in order to reduce GHG emissions (or energy consumption) on this site? Have you bought any new equipment, or have you changed the way you produce?” We document that more than 30% of managers report adopting optimization processes targeted at heating, waste heat recovery, industry-specific processes or machinery, and lighting.³²

Firms participating in the EU ETS were more likely to report making investments to optimize the use of process heat,³³ and to optimize processes specific to their industry, than non-ETS firms. We note that these correlations are descriptive and do not necessarily represent causal relationships. Nevertheless, in combination with our main results, these qualitative insights provide supporting evidence for the hypothesis that firms invested in new processes to reduce emissions.

Collectively, 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.

5.3. *Productivity*

What remains unresolved is that firms reduced emissions without any detectable contraction in economic activity, despite the fact that carbon pricing increased input costs. One hypothesis is that the ETS induced firms to make investments that increased productivity, offsetting any costs to the firm. However, the conditions under which such an interpretation can be rationalized are unclear. To explore this conjecture, we present a model of firm production that guides our evaluation. We use the structure of this model to estimate revenue-based total factor productivity (TFPR) and evaluate the effect of the EU ETS on measured . We estimate that, on average, the EU ETS had a positive but statistically insignificant effect on measured TFPR. We explore the implications of this finding first through the lens of a parsimonious baseline model. In contrast to our empirical findings, this baseline model predicts contractions in economic activity and weakly decreasing effects on productivity under the EU ETS. Next, we present a simple extension to the model incorporating the possibility of an alternative production technology that reduces the emissions intensity of production but requires paying a fixed switching cost. This extension rationalizes our empirical findings, delivering the possibility of weakly increasing effects on value added, employment, total capital, and productivity. The remainder of this section provides more details on each step of this investigation.

32. More than 15% of managers reported switching to natural gas, modernizing the compressed air system, innovating in the production processes, upgrading the energy management system, and also improving waste management and running employee awareness campaigns to reduce energy use.

33. As highlighted by [Ammar *et al.* \(2012\)](#), [Barma *et al.* \(2017\)](#), and [Chowdhury *et al.* \(2018\)](#), there is a sizable potential for waste recovery in the industrial sector. We thank an anonymous referee for pointing us to those studies.

5.3.1. Model environment. Consider a firm that uses capital K , energy services E , intermediate inputs M , and labour L to produce output Q .³⁴ Using energy services gives rise to CO₂ emissions when those services are produced with fossil fuels. We assume a Cobb-Douglas production function

$$Q = AE^{\alpha_E} K^{\alpha_K} M^{\alpha_M} L^{\alpha_L} \quad (5.1)$$

with returns-to-scale parameter $\gamma \equiv \alpha_K + \alpha_E + \alpha_M + \alpha_L$. The firm maximizes profits subject to an inverse isoelastic demand function³⁵

$$P = \Lambda^{\frac{1}{\mu}} Q^{\frac{1-\mu}{\mu}}, \quad (5.2)$$

where μ is the markup of prices over marginal costs and Λ is a demand shifter. Taking input prices, W_X , as given, a monopolistic firm's profit maximization problem becomes

$$V = \max_{E,K,M,L} \left\{ \Lambda^{\frac{1}{\mu}} Q^{\frac{1}{\mu}} - \sum_{X \in \{E,K,M,L\}} W_X X \right\}.$$

This leads to the following first-order conditions (FOCs):

$$X = \frac{\alpha_X}{\mu W_X} Q^{\frac{1}{\mu}} \Lambda^{\frac{1}{\mu}} \quad (5.3)$$

for $X \in \{E, K, M, L\}$, where we assume that all input factors are flexible. Using the production function, we solve for optimal output

$$Q^* = \left[A \prod_{X \in \{E,K,M,L\}} \left(\frac{\alpha_X}{W_X} \right)^{\alpha_X} \frac{\Lambda^{\frac{\gamma}{\mu}}}{\mu^{\frac{\gamma}{\mu}}} \right]^{\frac{\mu}{(\mu-\gamma)}}. \quad (5.4)$$

5.3.2. Measuring TFP. Like most studies on firm-level productivity, we do not observe physical output in our data. Instead, we observe revenue. Given the log-linear demand model function (4), we can write revenue as

$$R = \Lambda^{\frac{1}{\mu}} Q^{\frac{1}{\mu}}. \quad (5.5)$$

Under our assumptions about the production function, this can be restated as

$$R = \Lambda^{\frac{1}{\mu}} A^{\frac{1}{\mu}} \prod_X X^{\frac{\alpha_X}{\mu}}. \quad (5.6)$$

34. Following the literature, we use the term “energy services” to draw a distinction between the usage firms or households derive commonly from specific fuels (*e.g.* heating), and units of the fuel (*e.g.* tonnes of coal). We abstract from the fact that some energy services are derived from fuels that are not directly regulated at the firm level (*e.g.* electricity). We also do not explicitly model that some fossil fuels (*e.g.* natural gas in the chemical industry) are used as direct inputs in the production process rather than to derive energy services.

35. This demand function would follow from a Constant Elasticity of Substitution (CES) utility function in a monopolistic competition setting.

Taking natural logarithms and using lower-case letters to denote logged variables yields

$$r = \frac{1}{\mu}(\lambda + a) + \sum_x \frac{\alpha_x}{\mu} x. \quad (5.7)$$

From the FOC, we get the expression

$$X = \frac{1}{\mu} \frac{\alpha_x}{W_x} R \quad \forall X,$$

which we can rearrange to get

$$s_x \equiv \frac{W_x X}{R} = \frac{\alpha_x}{\mu}. \quad (5.8)$$

Substituting this into [equation \(9\)](#), Supplementary Material, yields

$$r = \frac{1}{\mu}(a + \lambda) + \sum_{x \neq K} s_x x + \left(\frac{\gamma}{\mu} - \sum_{x \neq K} \frac{\alpha_x}{\mu} \right) k, \quad (5.9)$$

where we have used the definition of the scale parameter $\gamma = \sum_x \alpha_x$. Rearranging terms leaves us with the following expression, which clarifies the notion of revenue productivity, as a composite of the technical efficiency a and the demand shifter λ :

$$\frac{1}{\mu}(a + \lambda) = r - \sum_{x \in \{e, l, m\}} s_x (x - k) - \frac{\gamma}{\mu} k. \quad (5.10)$$

We examine two measures of TFP that build on this formula.

The index number-based TFP residual. Consider

$$\tilde{\omega}_{it} = r_{it} - \sum_{x \in \{e, l, m\}} \check{s}_x (x_{it} - k_{it}) - k_{it}, \quad (5.11)$$

where \check{s}_x are the median expenditure shares for factors energy (E), intermediates (M), and labour (L). Subscript i indexes a particular firm and t a time period.

We use the median factor shares observed in the cross section of firms to reduce the impact of outliers. If firms flexibly adjust labour, intermediates, and energy (but not necessarily capital), then the productivity index in [equation \(13\)](#), Supplementary Material, represents a composite of the production function and demand shift parameters which can be interpreted as revenue productivity, $\tilde{\omega}_{it} \approx a_{it} + \lambda_{it}$, provided that the returns-to-scale and markup parameters γ and μ are close to one.

The estimation-based TFP residual. If firms have non-constant returns-to-scale γ and/or markups $\mu > 1$, then the above approach is unlikely to provide a consistent estimate of revenue productivity. In this case, we need an estimate of γ/μ to recover an index $\omega_{it} = (1/\mu)(a_{it} + \lambda_{it})$ of revenue productivity. This requires timing assumptions for ω_{it} and k_{it} . We assume an AR(1) process for ω_{it} :

$$\omega_{it} = \rho \omega_{it-1} + \eta_{it} \quad (5.12)$$

and that k_{it} is pre-determined in period t .³⁶ Under these assumptions we write:

$$\Theta_{it} - \frac{\gamma}{\mu} k_{it} = \rho \left(\Theta_{it-1} - \frac{\gamma}{\mu} k_{it-1} \right) + \eta_{it},$$

where $\Theta_{it} = r_{it} - \sum_{x \in \{e,l,m\}} \tilde{s}_x(x_{it} - k_{it})$. Rearranging yields a regression equation

$$\Theta_{it} = \frac{\gamma}{\mu} k_{it} + \rho \frac{\gamma}{\mu} k_{it-1} + \rho \Theta_{it-1} + \eta_{it} \quad (5.13)$$

that we estimate by OLS and compute revenue productivity as

$$\hat{\omega}_{it} = \Theta_{it} - \left(\frac{\hat{\gamma}}{\mu} \right) k_{it}. \quad (5.14)$$

In our empirical analysis, we focus on this measure of TFPR because it is less restrictive. Results are robust to using the index-based measure $\tilde{\omega}_{it}$.

5.3.3. The productivity effects of the EU ETS. This subsection shows that the predictions of the above model match some—but not all—of our empirical findings. This provides the motivation for extensions of the standard model that help to fully rationalize our empirical results, which we introduce in the next subsection.

In line with the literature (Baumol and Oates, 1988), we consider that the main effect of the EU ETS is to increase the price of energy services. If, as we have assumed above, profit-maximizing firms take factor costs as given, an increase in the price of carbon has no effect on a TFPR measure based on equation (12), Supplementary Material. As shown in Appendix Section D, Supplementary Material, TFPR remains equal to $\omega_{it} = (1/\mu)(a_{it} + \lambda_{it})$.

Contrary to this, Greenstone *et al.* (2012) model that environmental regulation reduces TFP, based on the notion that firms divert some exogenous share of their observed inputs to uses that do not contribute to observed output but that are needed to comply with the regulation. In the case of the EU ETS, such unproductive labour inputs may include employees that are in charge of measuring emissions, managing the permit holdings, and communicating with the regulator. In the context of our model, this would imply that the amount of effective labour is a fraction ν of total employment, *i.e.*

$$Q_{\text{ETS}} = A E^{\alpha_E} K^{\alpha_K} M^{\alpha_M} (\nu L)^{\alpha_L} = \nu^{\alpha_L} Q. \quad (5.15)$$

The FOCs are unchanged and therefore $Q_{\text{ETS}}^* = \nu^{\alpha_L} Q^*$. The effect of the EU ETS on TFP (ω) becomes

$$\Delta \omega = \omega_{\text{ETS}} - \omega = \frac{\partial \omega}{\partial q} \frac{\partial q}{\partial \ln \nu} \Delta \nu = \left(1 - \frac{\gamma}{\mu} \right) \frac{1}{\mu} \alpha_L \ln \nu,$$

which is negative since $\nu < 1$.

In column 8 of Table 3, we estimate that the EU ETS has no effect on measured TFPR, which is more consistent with our baseline model than the extension by Greenstone *et al.* (2012). However, we have other results that do not match the predictions of our baseline model, in

36. This is a simplified version of approaches further discussed in Forlani *et al.* (2023) and Aghion *et al.* (2023) and in line with similar approaches in the literatures; see, *e.g.* Klette and Griliches (1996), Olley and Pakes (1996), De Loecker and Warzynski (2012), and Akerberg *et al.* (2015).

particular the predicted negative effects on value added, employment, and capital. In our setting, value added is equal to

$$VA = R - W_E E - W_M M. \quad (5.16)$$

Hence,

$$\frac{\partial VA}{\partial W_E} = \frac{\partial R}{\partial W_E} - E - W_E \frac{\partial E}{\partial W_E} - W_M \frac{\partial M}{\partial W_M}.$$

Note from [equation \(5\)](#), Supplementary Material, that

$$\frac{\partial E}{\partial W_E} = -\frac{E}{W_E} + \frac{\alpha_E}{\mu W_E} \frac{\partial R}{\partial W_E}$$

and

$$\frac{\partial M}{\partial W_E} = \frac{\alpha_M}{\mu W_E} \frac{\partial R}{\partial W_E}$$

so that

$$\frac{\partial VA}{\partial W_E} = \frac{\partial R}{\partial W_E} \left(1 - \frac{\alpha_E}{\mu}\right) = \frac{1}{\mu} \frac{\partial Q}{\partial W_E} \left(1 - \frac{\alpha_E}{\mu} - \frac{\alpha_M}{\mu}\right) < 0$$

as $\partial Q/\partial W_E < 0$ and $\alpha_E + \alpha_M < \mu$. A higher cost from carbon pricing implies that firms should reduce their output and, as a consequence, factor demand for all inputs. *Ceteris paribus*, value added, capital, and employment should all fall. Instead, we estimate a significant increase in capital and positive (but statistically insignificant) effects of the ETS on value added, employment, and measured TFP. To rationalize those results, we need to augment the baseline model.

5.3.4. A model with technology switching. We show that our empirical results are consistent with a model where firms can respond to higher carbon prices by switching to an alternative production technology that saves energy thereby reducing CO₂ emissions. We also assume that this technology is characterized by higher TFP. Why would firms not adopt this technology absent carbon pricing? We assume that adoption requires firms to pay a fixed switching cost, κ . The resulting trade-off between higher up-front investments and lower running costs is a common feature of many clean technologies. For example, combined heat and power generation or waste heat recovery technologies typically require a re-organization of production facilities alongside up-front investments in additional equipment which lead to subsequent reductions in running costs. Consistent with this narrative, these technologies featured prominently among the production changes that managers at ETS firms reported in interview data discussed in [Section 5.2](#) (cf. [Table C.1](#), Supplementary Material).³⁷

Formally, we assume that the alternative (clean) technology state is characterized by

$$\alpha'_E = \alpha_E - \xi_\alpha, \quad (5.17)$$

$$\alpha'_K = \alpha_K + \xi_\alpha, \quad (5.18)$$

$$A' = A + \xi_A. \quad (5.19)$$

37. Firms may also innovate to reduce emissions ([Calel and Dechezleprêtre, 2016](#)), another fixed-cost investment likely to boost TFP.

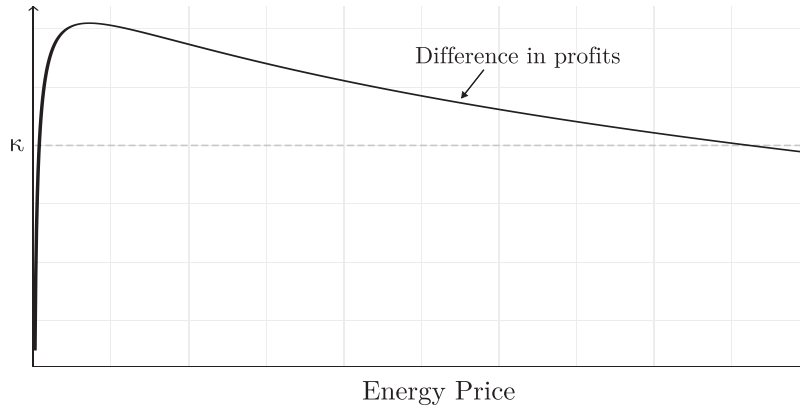


FIGURE 4
Relative profits from technology switching

That is this alternative technology is less energy intensive and more capital-intensive (by ξ_a), and has a higher TFP (by ξ_A). Firms apply a discount rate of r and will therefore switch to the new technology if the present discounted value of doing so exceeds the switching cost of κ , *i.e.*

$$(\Pi' - \Pi) \frac{1+r}{r} > \kappa, \quad (5.20)$$

where Π' and Π denote per period profits of the firm in the “clean” and “dirty” technology states, respectively. We argue that, prior to the introduction of the ETS, the marginal firm may not have been willing to make the fixed-cost investment. However, following the introduction of a carbon price, which increases energy prices and the cost of using the traditional technology, the present discounted value of making that investment may exceed the fixed cost of switching technologies.³⁸

Figure 4 visualizes a parameterization of the difference in profits between the “clean” and “dirty” technology states for a range of energy prices. The relationship has a hyperbolic inverted-U shape, which tends to minus infinity if the energy price tends to zero, and tends to zero as the energy price tends to infinity. This is because using a more energy intensive production technology is (infinitely) more profitable if energy costs nothing and, at the other extreme, because any use of energy makes production unprofitable if energy is infinitely costly. In between, there is a range of energy prices where the extra discounted profit from adopting exceeds the switching cost κ —provided κ is not too high—leading the firm to adopt the clean production technology.

The introduction of carbon pricing results in two scenarios. If the carbon price is too high or too low, no technology switching occurs. Firms remain with the traditional production technology. Higher marginal costs reduce energy usage (and hence emissions) as well as value added and other inputs—firms contract. As discussed in the previous section, measured TFP remains unchanged. We provide a parameterization of this scenario in Figure 5, panel a).

38. In practice, managers may use simpler decision rules, such as maximum payback time, which amounts to using high discount rates in equation (22), Supplementary Material. In interviews with managers of French manufacturing firms (further described in Appendix C.2, Supplementary Material), Martin *et al.* (2014b) asked about the maximum payback time required for an energy efficiency enhancing measure the firm had considered but not adopted. The median (mean) answer was 3 (3.6) years. Without carbon pricing, many energy efficiency investments may not pay back the investment cost fast enough to satisfy this criterion.

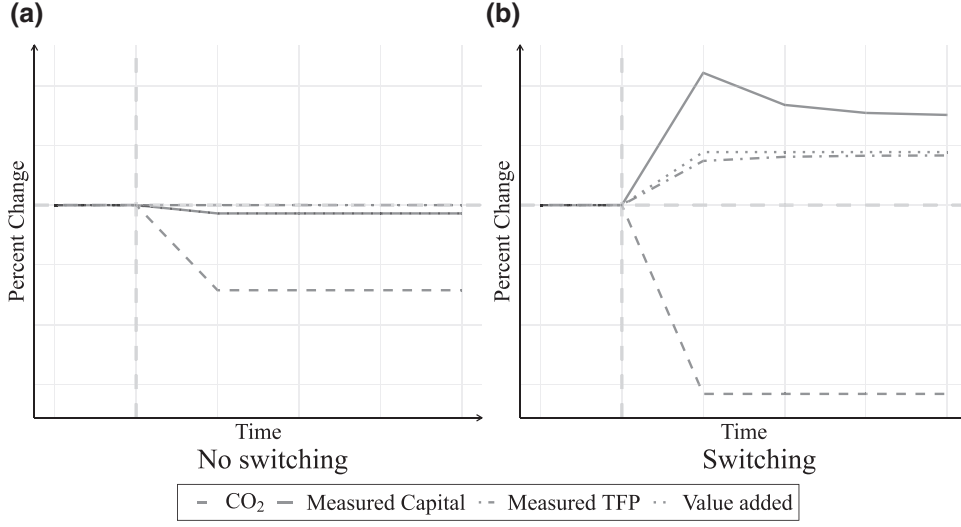


FIGURE 5

The effect of carbon pricing with and without technology switching

Notes: The figure illustrates the dynamics in outcome variables for different technology states. With technology switching (b), CO₂ emissions fall more sharply than without switching (a). The productivity-enhancing effect of the new technology leads to an increase in value added. Measured TFP trails the increase in actual TFP because measured capital overstates the amount of productive capital. Since data on one-off investments to switch production technologies is not separately available from other investments in fixed capital, measured capital likely includes any switching costs. Measured capital exceeds pre-policy levels and subsequently depreciates geometrically, reducing the bias in measured TFP.

Within the intermediate range of carbon prices, technology switching occurs. After switching, output—along with value added—could increase or decrease compared to a state without carbon pricing. To show which factors determine the direction of this change, we write log output as

$$q(\zeta_a, \zeta_A) = \frac{\mu}{\mu - \gamma} [a + \zeta_A + (\alpha_E - \zeta_a) (\ln(\alpha_E - \zeta_a) - \ln W_E) + (\alpha_K + \zeta_a) (\ln(\alpha_K + \zeta_a) - \ln W_K) + \sum_{X \in \{L, M\}} [\alpha_X (\ln \alpha_X - \ln W_X)]],$$

where $a \equiv \ln A$. A firm that was initially using the traditional technology ($\zeta_a = \zeta_A = 0$) and switches technologies due to an increase in the carbon price will see its output affected via changes in ζ_a , ζ_A , and $\ln W_E$:

$$dq(0, 0) = \frac{\partial q(\zeta_a = 0, \zeta_A = 0)}{\partial \ln W_E} d \ln W_E + \frac{\partial q(\zeta_a = 0, \zeta_A = 0)}{\partial \zeta_a} d \zeta_a + \frac{\partial q(\zeta_a = 0, \zeta_A = 0)}{\partial \zeta_A} d \zeta_A.$$

The first term captures the direct effect of the carbon price on output. It is strictly negative because it captures an increase in marginal costs:

$$\frac{\partial q(0, 0)}{\partial \ln W_E} = -\frac{\mu}{\mu - \gamma} \alpha_E < 0.$$

The second term captures the effect on output resulting from the reduction in energy intensity due to the new technology:

$$\frac{\partial q(0, 0)}{\partial \zeta_\alpha} = \frac{\mu}{\mu - \alpha_E - \alpha_K} [-(\ln \alpha_E - \ln W_E) + (\ln \alpha_K - \ln W_K)]. \quad (5.21)$$

The sign of this term is ambiguous. The change in technology reduces the energy intensity of production and increases the intensity of capital. This lowers marginal costs if energy is expensive, since the new technology relies on it less. If, however, energy were cheap compared to capital before the arrival of carbon pricing, then the transition to using capital more intensely may increase marginal costs. Likewise, the relative size of α_E and α_K matters. When α_E is low relative to α_K , the second term is more likely to be positive. Intuitively, this is because for every inframarginal unit of energy that becomes less effective when the firm switches technologies, there is more than one inframarginal unit of capital that becomes more effective.

The third term captures the output effect of an increase in TFP by ζ_α , which is unambiguously positive:

$$\frac{\partial q(0, 0)}{\partial \zeta_A} = \frac{\mu}{\mu - \gamma} > 0.$$

The overall effect on output and, consequently, value added and profits depends on the relative magnitude of these terms. Panel (b) of Figure 5 illustrates a parameterization of the model in which technology switching induces a reduction in CO₂, a net increase in output (and value added) as well as the corresponding increases in measured capital and measured TFP that coincide with switching.³⁹

Our empirical results, which document reductions in emissions, increases in measured capital, and weakly increasing effects on value added and measured TFPR, are consistent with the case in which technology switching induces increases in TFP and reductions in marginal cost sufficient enough to offset the contractionary effects of carbon pricing.

6. AGGREGATE CARBON SAVINGS

We combine our estimates with the aggregated microdata on CO₂ emissions to gauge the extent to which the EU ETS might have driven aggregate emission reductions since 2005. Details on the calculations below can be found in Appendix D, Supplementary Material.

The black line in Figure 6, constructed using our microdata, depicts observed aggregate industrial CO₂ emissions in France between 1996 and 2012. We observe that aggregate emissions have been falling over time, and that the decline has been steeper in recent years.

We see a substantial aggregate drop in emissions starting in 2005 at the start of the ETS and again in 2008 at the start of phase II. The dashed line plots emissions in 2004 as a benchmark. These findings are consistent with our empirical evidence, but the question remains: how much did the ETS contribute to these aggregate reductions?

We calculate that between 2005 and 2012 aggregate emissions would on average have been 5.4 million tonnes higher each year if there were no EU ETS. Compared to 2004 emissions, this accounts for 28% of the aggregate emissions reduction during this period. Using the linear trend in emissions prior to 2005 as a benchmark instead of emissions in 2004 would lead us to attribute

39. The purpose of this calibration is to illustrate the possible range of outcomes. We leave more substantive calibration exercises for future research.

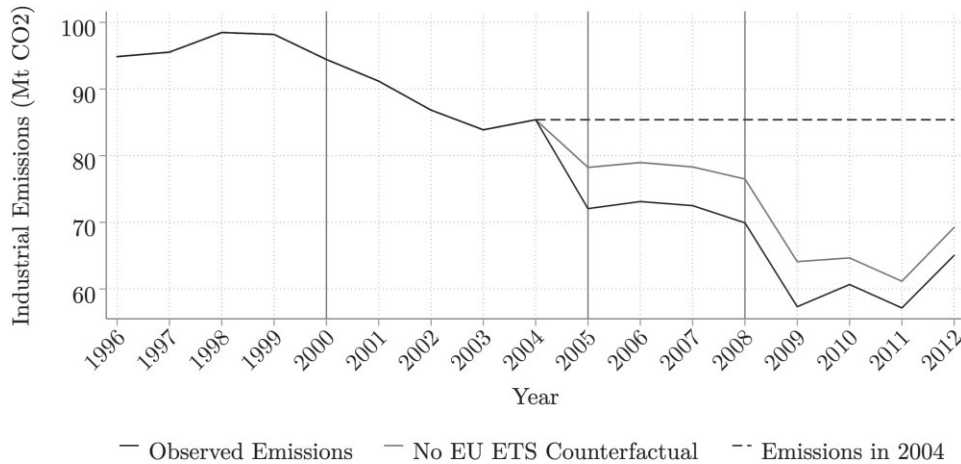


FIGURE 6

The effect of the EU ETS on aggregate emissions reductions

Notes: The black line presents the aggregate time series for industrial emissions in France, measured in millions of tonnes of CO₂. The dark grey 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 dashed black line represents the level of emissions in 2004 as a benchmark. Source: Authors calculations based on French microdata and Eurostat data.

47% of the aggregate emissions reduction during this period to the EU ETS. These calculations highlight the importance of causal research designs for evaluating the efficacy of climate policy. Of the aggregate emissions reductions in our data, 53–72% are driven by other factors, such as structural economic change, energy efficiency improvements, or the Great Recession. Drawing inferences about the effectiveness of the ETS based on aggregate patterns and trend-breaks would lead us to vastly overestimate the efficacy of the EU ETS.

The emissions reductions observed in the data occurred in spite of carbon prices averaging a relatively low \$21.35 per tonne (\$2017) during phase II. The average abatement costs per tonne of CO₂ must have been lower, since it would otherwise 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 favourably with that of non-market-based approaches (Carlson *et al.*, 2000; Fowlie *et al.*, 2012; Gillingham and Stock, 2018). In Figure 7, we compare the estimated cost per tonne of CO₂ (\$2017) for twenty-five climate change mitigation policies. The estimate for the EU ETS is based on the maximum price during phase II, which was \$52.68. This is an upper bound cost estimate—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 and Stock (2018). Even when we use the maximum cost per tonne of CO₂, the EU ETS is ranked seventh. If we use the average phase II price instead (\$21.35), which is still likely to be very conservative, the EU ETS is ranked fifth. We caveat that this exercise assumes that the EUA price is unaffected by the other energy and climate policies discussed in Appendix B.4, Supplementary Material. While we do not think that these policies differentially affected ETS firms, their existence may have had an aggregate effect, resulting in a lower equilibrium permit price. This would have the effect of making the ETS as a whole (*i.e.* including the electricity sector) appear cheaper than it would have been if these policies did not exist.

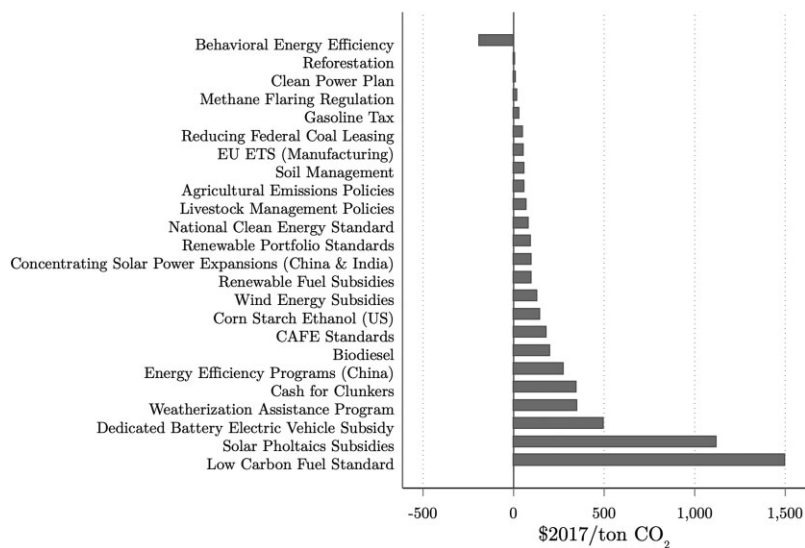


FIGURE 7
Comparing the EU ETS to other climate change mitigation policies

Notes: This figure ranks different climate change policies by the estimated cost of reducing a tonne 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 1 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 tonne of CO₂ was \$52.68. The actual cost was likely far lower, since 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 seventh out of twenty-five. The cost of other policies are taken from [Gillingham and Stock \(2018\)](#). Where multiple estimates exist for the same policy, we take the average across all estimates.

7. 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. We find little evidence that carbon leakage played a meaningful role in contributing to emissions reductions, indicating that, at least in this context, the EU ETS helped to mitigate global climate change. Our findings are consistent with firms paying an up-front fixed cost to invest in alternative “clean” production technologies that reduce marginal variable costs. The results suggest that, when firms make such investments, decarbonization may only be costly in the transition phase, rather than in the long term.

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 they may deliver improvements in environmental quality, such non-market-based policies have been shown to negatively affect firm performance ([Becker and Henderson, 2000](#); [Greenstone, 2002](#); [Greenstone et al., 2012](#); [Walker, 2013](#); [He et al., 2020](#)).

We note caveats. First, despite the significant effect that the EU ETS has had on emissions, these results should not be taken as a blanket 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—emissions trading systems. Our context is one in which compliance is high and corruption low. Second, while we do not estimate any significant contractions in economic activity, this does not imply that emissions reductions were made without cost. 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.

Acknowledgments. We thank Bård Harstad for his comments and guidance, and three anonymous referees whose suggestions have significantly improved the paper. We are grateful to Michael Best, Meredith Fowlie, Andreas Gerster, Michael Greenstone, Stephen Holland, Xavier Koch, Jakob Lehr, Mar Reguant, Johannes Spinnewijn, and John Van Reenen for helpful thoughts and discussions. We are also grateful to seminar participants at Harvard, 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, the National Bank of Belgium, European Union's Horizon 2020 Research and Innovation Programme under grant agreement nos. 308481 and 865181, the German Research Foundation (DFG) through CRC TR 224 (Project B07) 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" programme (reference: ANR-10-EQPX-17—Centre d'accès sécurisé distant aux données—CASD). All errors and omissions are our own.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Data Availability Statement

The data used in this paper are the property of the French government. Details on the process for applying to access the data, and replication code to reproduce our results, are available in Zenodo at <https://doi.org/10.5281/zenodo.10791032>.

REFERENCES

- ABADIE, A. and SPIESS, J. (2022), "Robust Post-Matching Inference", *Journal of the American Statistical Association*, **117**, 983–995.
- ACKERBERG, D., CAVES, K. and FRAZER, G. (2015), "Identification Properties of Recent Production Function Estimators", *Econometrica*, **83**, 2411–2451.
- AGHION, P., BENABOU, R., MARTIN, R., *et al.* (2023), "Environmental Preferences and Technological Choices: Is Market Competition Clean or Dirty?" *American Economic Review: Insights*, **5**, 1–20.
- AMMAR, Y., JOYCE, S., NORMAN, R., *et al.* (2012), "Low Grade Thermal Energy Sources and Uses from the Process Industry in the UK", *Applied Energy*, **89**, 3–20.
- ANDERSEN, B. and DI MARIA, C. (2011), "Abatement and Allocation in the Pilot Phase of the EU ETS", *Environmental and Resource Economics*, **48**, 83–103.
- BARMA, M., SAIDUR, R., RAHMAN, S., *et al.* (2017), "A Review on Boilers Energy Use, Energy Savings, and Emissions Reductions", *Renewable and Sustainable Energy Reviews*, **79**, 970–983.
- BAUMOL, W. (1972), "On Taxation and the Control of Externalities", *American Economic Review*, **62**, 307–322.
- BAUMOL, W. and OATES, W. (1971), "The Use of Standards and Prices for Protection of the Environment", *Swedish Journal of Economics*, **73**, 42–54.
- BAUMOL, W. and OATES, W. (1988), *The Theory of Environmental Policy* (Cambridge: Cambridge University Press).
- BECKER, R. and HENDERSON, J. (2000), "Effects of Air Quality Regulations on Polluting Industries", *Journal of Political Economy*, **108**, 379–421.
- BERGOUNHON, F., LENOIR, C. and MEJEAN, I. (2018), "A Guideline to French Firm Level Trade Data", Unpublished, CREST-Ecole Polytechnique.
- BORGHESI, S., FRANCO, C. and 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**, 219–256.
- BURKE, M., CRAXTON, M., KOLSTAD, C., *et al.* (2016), "Opportunities for Advances in Climate Change Economics", *Science*, **352**, 292–293.
- CALEL, R. (2020), "Adopt or Innovate: Understanding Technological Choices Under Cap-and-Trade", *American Economic Journal: Economic Policy*, **12**, 170–201.
- CALEL, R. and DECHEZLEPRÊTRE, A. (2016), "Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market", *Review of Economics and Statistics*, **98**, 173–191.
- CARLSON, C., BURTRAW, D., CROPPER, M., *et al.* (2000), "Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade?" *Journal of Political Economy*, **108**, 1292–1326.
- CHOWDHURY, J. I., HU, Y., HALTAS, I., *et al.* (2018), "Reducing Industrial Energy Demand in the UK: A Review of Energy Efficiency Technologies and Energy Saving Potential in Selected Sectors", *Renewable and Sustainable Energy Reviews*, **94**, 1153–1178.
- DECHEZLEPRÊTRE, A., NACHTIGALL, D. and VENMANS, F. (2023), "The Joint Impact of the European Union Emissions Trading System on Carbon Emissions and Economic Performance", *Journal of Environmental Economics and Management*, **118**, 102758.

- DE JONGHE, O., MULIER, K. and SCHEPENS, G. (2020), "Going Green by Putting a Price on Pollution: Firm-Level Evidence from the EU" (Working Papers 390, National Bank of Belgium).
- DE LOECKER, J. and WARZYNSKI, F. (2012), "Markups and Firm-Level Export Status", *American Economic Review*, **102**, 2437–2471.
- EATON, J., KORTUM, S. and KRAMARZ, F. (2011), "An Anatomy of International Trade: Evidence from French Firms", *Econometrica*, **79**, 1453–1498.
- EGENHOFER, C., ALESSI, M., FUJIWARA, N., *et al.* (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. and 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**, 267–287.
- ELLERMAN, A. D., CONVERY, F. J. and DE PERTHUIS, C. (2010), *Pricing Carbon: The European Union Emissions Trading Scheme* (Cambridge: Cambridge University Press).
- ELLERMAN, A. D. and FEILHAUER, S. M. (2008), "A Top-Down and Bottom-Up Look at Emissions Abatement in Germany in Response to the EU ETS" (Working Paper 08-017, Center for Energy and Environmental Policy Research, Massachusetts Institute of Technology).
- ELLERMAN, D., MARCANTONINI, C. and ZAKLAN, A. (2016), "The European Union Emissions Trading System: Ten Years and Counting", *Review of Environmental Economics and Policy*, **10**, 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. and REGUANT, M. (2014), "Pass-Through of Emissions Costs in Electricity Markets", *American Economic Review*, **104**, 2872–2899.
- FORLANI, E., MARTIN, R., MION, G., *et al.* (2023), "Unraveling Firms: Demand, Productivity and Markups Heterogeneity", *The Economic Journal*, **133**, 2251–2302.
- FOWLIE, M., HOLLAND, S. and MANSUR, E. (2012), "What Do Emissions Markets Deliver and to Whom? Evidence from Southern California's NOx Trading Program", *American Economic Review*, **102**, 965–993.
- FOWLIE, M., REGUANT, M. and RYAN, S. (2016), "Market-Based Emissions Regulation and Industry Dynamics", *Journal of Political Economy*, **124**, 249–302.
- FREYALDENHOVEN, S., HANSEN, C. and SHAPIRO, J. (2019), "Pre-event Trends in the Panel Event-Study Design", *American Economic Review*, **109**, 3307–3338.
- GERSTER, A., LEHR, J., PIEPER, S., *et al.* (2021), "The Impact of Carbon Trading on Industry: Evidence from German Manufacturing Firms" (Mimeo, University of Mannheim).
- GILLINGHAM, K. and STOCK, J. (2018), "The Cost of Reducing Greenhouse Gas Emissions", *Journal of Economic Perspectives*, **32**, 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**, 1175–1219.
- GREENSTONE, M., LIST, J. and SYVERSON, C. (2012), "The Effects of Environmental Regulation on the Competitiveness of US Manufacturing" (Working Paper No. 18392, NBER).
- HAHN, R. W. (1989), "Economic Prescriptions for Environmental Problems: How the Patient Followed the Doctor's Orders", *Journal of Economic Perspectives*, **3**, 95–114.
- HE, G., WANG, S. and ZHANG, B. (2020), "Watering Down Environmental Regulation in China", *Quarterly Journal of Economics*, **135**, 2135–2185.
- HECKMAN, J., ICHIMURA, H. and TODD, P. (1997), "Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme", *Review of Economic Studies*, **64**, 605–654.
- HECKMAN, J., ICHIMURA, H. and TODD, P. (1998), "Matching as an Econometric Evaluation Estimator", *Review of Economic Studies*, **65**, 261–294.
- HENDREN, N. and FINKELSTEIN, A. (2020), "Welfare Analysis Meets Causal Inference", *Journal of Economic Perspectives*, **34**, 146–167.
- HENDREN, N. and SPRUNG-KEYSER, B. (2020), "A Unified Welfare Analysis of Government Policies", *Quarterly Journal of Economics*, **135**, 1209–1318.
- HINTERMANN, B. (2016), "Pass-Through of CO₂ Emission Costs to Hourly Electricity Prices in Germany", *Journal of the Association of Environmental and Resource Economists*, **3**, 857–891.
- HO, D., IMAI, K., KING, G., *et al.* (2007), "Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference", *Political Analysis*, **15**, 199–236.
- JARAITE, J. and DI MARIA, C. (2016), "Did the EU ETS Make a Difference? An Empirical Assessment Using Lithuanian Firm-Level Data", *Energy Journal*, **37**, 1–23.
- KLEMETSEN, M. E., ROSENDAHL, K. E. and JAKOBSEN, A. L. (2020), "The Impacts of the EU ETS on Norwegian Plants' Environmental and Economic Performance", *Climate Change Economics*, **11**, 2050006.
- KLETTE, T. and GRILICHES, Z. (1996), "The Inconsistency of Common Scale Estimators When Output Prices Are Unobserved and Endogenous", *Journal of Applied Econometrics*, **11**, 343–61.
- KOCH, N. and BASSE MAMA, H. (2019), "Does the EU Emissions Trading System Induce Investment Leakage? Evidence from German Multinational Firms", *Energy Economics*, **81**, 479–492.

- LÖSCHEL, A., LUTZ, B. J. and MANAGI, S. (2019), "The Impacts of the EU ETS on Efficiency and Economic Performance. An Empirical Analysis for German Manufacturing Firms", *Resource and Energy Economics*, **56**, 71–95.
- MARIN, G., MARINO, M. and 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. and WAGNER, U. (2014a), "The Impact of a Carbon Tax on Manufacturing: Evidence from Microdata", *Journal of Public Economics*, **117**, 1–14.
- MARTIN, R., MUÜLS, M., DE PREUX, L., *et al.* (2014b), "Industry Compensation Under Relocation Risk: A Firm-Level Analysis of the EU Emissions Trading Scheme", *American Economic Review*, **104**, 1–24.
- MARTIN, R., MUÜLS, M. and WAGNER, U. J. (2016), "The Impact of the European Union Emissions Trading Scheme on Regulated Firms: What Is the Evidence After Ten Years?" *Review of Environmental Economics and Policy*, **10**, 129–148.
- MAYER, T., MELITZ, M. and OTTAVIANO, G. (2014), "Market Size, Competition, and the Product Mix of Exporters", *American Economic Review*, **104**, 495–536.
- MONTGOMERY, W. (1972), "Markets in Licenses and Efficient Pollution Control Programs", *Journal of Economic Theory*, **5**, 395–418.
- NORDHAUS, W. (1977), "Economic Growth and Climate: The Carbon Dioxide Problem", *American Economic Review*, **67**, 341–346.
- NORDHAUS, W. (2001), "Global Warming Economics", *Science*, **294**, 1283–1284.
- OLLEY, G. S. and PAKES, A. (1996), "The Dynamics of Productivity in the Telecommunications Equipment Industry", *Econometrica*, **64**, 1263–97.
- PIGOU, A. (1920), *The Economics of Welfare* (London: Macmillan).
- RAMBACHAN, A. and ROTH, J. (2023), "A More Credible Approach to Parallel Trends", *Review of Economic Studies*, **90**, 2555–2591.
- ROTH, J. (2022), "Pretest with Caution: Event-Study Estimates After Testing for Parallel Trends", *American Economic Review: Insights*, **4**, 305–322.
- RYAN, S. P. (2012), "The Costs of Environmental Regulation in a Concentrated Industry", *Econometrica*, **80**, 1019–1061.
- TIETENBERG, T. (1973), "Controlling Pollution by Price and Standard Systems: A General Equilibrium Analysis", *Swedish Journal of Economics*, **75**, 193–203.
- VOGT-SCHILB, A. and HALLEGATTE, S. (2014), "Marginal Abatement Cost Curves and the Optimal Timing of Mitigation Measures", *Energy Policy*, **66**, 645–653.
- WALKER, W. R. (2013), "The Transitional Costs of Sectoral Reallocation: Evidence from the Clean Air Act and the Workforce", *Quarterly Journal of Economics*, **128**, 1787–1835.
- ZAKLAN, A. (2023), "Coase and Cap-and-Trade: Evidence on the Independence Property from the European Electricity Sector", *American Economic Journal: Economic Policy*, **15**, 526–558.
- ZUBIZARRETA, J. R., PAREDES, R. D. and ROSENBAUM, P. R. (2014), "Matching for Balance, Pairing for Heterogeneity in an Observational Study of the Effectiveness of For-Profit and Not-For-Profit High Schools in Chile", *The Annals of Applied Statistics*, **8**, 204–231.