

# Emissions Trading, Firm Behavior, and the Environment: Evidence from French Manufacturing Firms\*

Jonathan Colmer

University of Virginia

Mirabelle Muûls

Imperial College London

Ralf Martin

Imperial College London

Ulrich J. Wagner

University of Mannheim

July 21, 2018

## Abstract

Market-based instruments of regulation hold the promise of minimizing total compliance costs by letting regulated firms choose *how* to comply. When regulation is incomplete, regulated firms have an additional margin of choice concerning the *degree* of compliance. Under these circumstances, opportunistic firm behavior not only minimizes compliance cost but it can also undermine the efficacy of regulation. We analyze this trade-off in the context of regulation aimed at internalizing global environmental externalities that cause climate change. Using administrative data on French manufacturing firms we establish that the EU Emissions Trading Scheme caused treated firms to reduce carbon dioxide emissions by 13.5% relative to untreated firms. This substantial abatement was achieved by switching into (zero-carbon) electricity consumption rather than by down-scaling production. Further, we analyze whether firms complied in ways that would not reduce emissions globally, namely by (i) shifting pollution from regulated to unregulated plants within firm or by (ii) increasing foreign-sourced intermediates to replace domestic production. Our results support the view that regulated firms used the former channel, but there is no evidence that regulated firms imported more intermediates from abroad in order to reduce their emissions. Finally, we show that regulated firms ramped up emissions following the announcement but prior to the beginning of the policy, possibly to increase rents from free permit allocation.

---

\*First Version: February 2013. Previous versions of this paper have circulated under the title “The Causal Effects of the European Union Emissions Trading Scheme: Evidence from French Manufacturing Plants”. Correspondence: m.muuls@imperial.ac.uk. Affiliations: Colmer – Department of Economics, University of Virginia, USA; Martin – Imperial College Business School, Imperial College London and Centre for Economic Performance, LSE, UK; Muûls – Imperial College Business School, Imperial College London, UK; Wagner – Department of Economics, University of Mannheim and ZEW, Germany. We thank Michael Best, Meredith Fowlie, Andreas Gerster, Michael Greenstone, Mar Reguant, Johannes Spinnewijn, and John Van Reenen for very helpful thoughts, comments, and discussions. We are also grateful to seminar participants at Cambridge, Geneva, LSE, Mannheim, Toulouse, at IZA and many 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). Research funding from the Economic and Social Research Council (grant ES/J006742/1), the ESRC Centre for Climate Change Economics and Policy, the DFG Collective Research Centre TR 224 and the Grantham Foundation is gratefully acknowledged. All errors and omissions are our own.

# 1 Introduction

The presence of market failure calls for government intervention. However, if government intervention is to be effective, one must consider both its intended and unintended consequences. If the behavioral responses of agents to corrective policies are not taken into consideration, then the social benefits sought by such interventions may be substantially eroded, and potentially more costly than in the absence of intervention.

One of the clearest examples of market failure in the 21st century is climate change. Indeed, the Intergovernmental Panel on Climate Change, in its 5th Assessment Report (IPCC, 2014) presents convincing evidence that climate change is driven by anthropogenic greenhouse gas (GHG) emissions that mix (nearly) uniformly in the global atmosphere. In view of the dramatic consequences that unchecked climate change is predicted to have on human livelihoods, ecosystems, and built infrastructures worldwide, it is clear that globally harmonized action on climate change is best suited to internalizing the externalities associated with GHG emissions.

In spite of this insight, the existing policy architecture for mitigating climate change is a patchwork of national and regional policies. Given the global integration of economic activity, this has raised concerns about “carbon leakage”, which undermines the effectiveness of unilateral policies. If firms subject to regulation of carbon emissions have the ability to outsource production to markets without such regulation then carbon emissions “leak” from regulated to unregulated jurisdictions. This leakage results in an absence of – or even negative – aggregate reduction in emissions. In such a scenario, unilateral environmental regulation achieves little in terms of environment protection while potentially hurting the competitiveness of regulated firms and “killing jobs”, as industry advocates never fail to point out in the debate about effective public policies for addressing global climate change. The degree to which such policies curb the externality that they seek to regulate is thus of first-order concern. One of the reasons that this debate over the consequences of environmental policy has carried on is the paucity of conclusive empirical evidence.

In this paper we shed light on these issues through an examination of the European Union (EU) Emissions Trading Scheme (ETS).<sup>1</sup> As the name suggests, the ETS is a cap-and-trade scheme for carbon dioxide (CO<sub>2</sub>) emissions. Regulated entities in the power and industrial sectors, among others, need a valid pollution permit for each ton of CO<sub>2</sub> they emit per year. They obtain permits from the government or by bidding in permit auctions. Subsequently, they can trade permits in an international permit market or hold on to them for future use. The total supply of emission rights is capped and bound to decrease over time. The

---

<sup>1</sup>For the sake of brevity, we refer to the policy as “the ETS” throughout this paper.

current level of the cap corresponds to approximately 45% of EU emissions and 5% of global emissions. This makes the ETS the world’s largest international system for trading GHG emissions, as well as the EU’s flagship climate policy.

To evaluate the aggregate environmental consequences of the ETS, we study how the policy affected firm behavior using rich, administrative data on French manufacturing firms between 1994 and 2012. Our identification strategy exploits variation in treatment status across firms which is induced by underlying variation in the number and capacity rating of energy-intensive installations which are targeted by the ETS participation rules. We argue that this variation is independent of the outcomes once we condition on differences in firm characteristics of regulated vs. unregulated firms. We accomplish this conditioning using semi-parametric matching techniques and subsequently estimate the average treatment effect on treated firms by performing differences-in-differences (DiD) estimations on the matched sample.

We find that during the second phase of the ETS (2008-2012) regulated firms reduced their emissions by 13.5%, relative to the year 2000, compared to unregulated firms. In addition, we find no statistically significant changes to employment, value added, or the capital stock, suggesting that the effects of the ETS on the competitiveness of regulated firms have been limited.

To arrive at a proper assessment of the cost effectiveness of the ETS, one needs to know whether the estimated emission reductions are real in a global sense, or offset by carbon leakage. To shed light on this, we evaluate the degree to which emission reductions may have arisen as a result of firm-level adjustments to the organization of the production process that shift emissions elsewhere instead of avoiding them. We propose that there are two means through which a reallocation from regulated to unregulated facilities may occur. First, firms may reallocate emissions from regulated to unregulated facilities within the firm. Second, they may adjust their supply chains, resulting in a reallocation of “dirty” production from regulated firms to unregulated firms. Because unregulated firms are likely to be located abroad, we explore the extent to which sourcing of intermediate inputs operates through international trade.

In support of the hypothesis that emissions are reallocated across plants within the firm, we find that firms with no untreated plants – firms that have limited scope for within-firm reallocation – reduce emissions during Phase II by 23%. By contrast, firms with both regulated and unregulated establishments – partially regulated firms – do not reduce their emissions at all. Further to this point, we use plant-level data to directly evaluate within-firm reallocations. The point estimates, while not statistically significant, suggest that regulated plants in partially-regulated firms reduced their emissions by 11%, with a corresponding 15%

increase in emissions for unregulated plants in partially-regulated firms. These findings suggest that aggregate emission reductions could have been substantially larger in the absence of within-firm reallocation.

In terms of the potential for between-firm reallocation, we find that the ETS had no effects on imports (neither total imports nor imports from outside the ETS market). In addition, regulated firms did not make any changes to the number of countries they imported from, or to the number of products that they imported. Collectively, these findings suggest that the 23% reduction in emissions for fully-regulated firms are likely real in a global sense.

The remainder of this paper is organized as follows. The next section provides the policy background for our analysis and reviews the related literature. Section 3 presents the various datasets we use and explains the linking procedures. Section 4 introduces the research design and Section 5 presents the results. Section 6 concludes.

## 2 Background

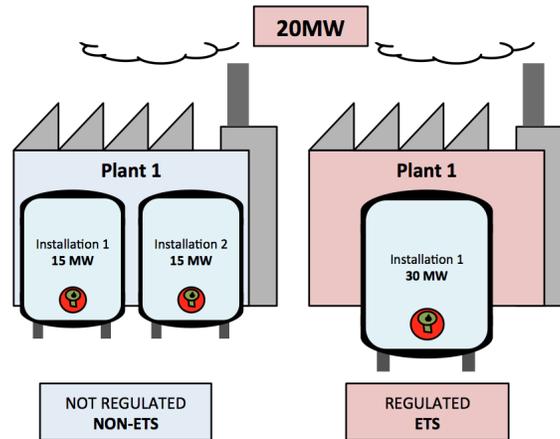
### 2.1 The European Union Emissions Trading Scheme

Since 2005, the EU Commission and member states have been running a large-scale cap-and-trade scheme with the aim to limit and reduce the CO<sub>2</sub> emissions of more than 11,000 power and manufacturing plants in 31 countries. The ETS is currently in its third phase, which will last until 2020. The first phase took place from 2005 to 2007, and the second one from 2008 to 2012. The fourth trading phase will start in 2021. The establishment of an ETS for the European Union was first evoked in the “Green Paper on greenhouse gas emissions trading within the European Union” (European Commission, 2000). To this publication followed a period of consultation and the submission of comments by businesses, NGOs and governments. The summary of these comments was published in May 2001 (European Commission, 2001), presenting more information and clarity to the shape that the ETS would be taking. Notwithstanding this, the Emissions Trading Directive 2003/87/EC was not published until 2003. In what follows, we take a conservative view and consider that an announcement effect of the ETS could have emerged as early as in 2001.

Participating installations are required to surrender, at the end of each year, one European Union Allowance (EUA) for each ton of gas they have emitted over the year. Participants are free to buy additional EUAs or sell excess EUAs on an international market. They can also bank or borrow EUAs to balance their needs across years. A comprehensive review of the history and structure of the ETS can be found in [Ellerman et al. \(2016\)](#).

A brief discussion of the participation criteria is needed in order to motivate how we

Figure 1: Regulation status at the plant level



distinguish between treatment and control firms. These criteria were first spelled out in the Emissions Trading Directive 2003/87/EC and transposed into national laws.<sup>2</sup> In France, participation in the ETS is mandatory for all combustion installations with a rated thermal input of 20 mega Watts (MW) or more. This order of magnitude is mostly relevant to heat and power generation, regardless of the industry. While combustion installations account for the lion's share of total emissions under the cap, industrial installations are also included if they (i) specialize in certain industrial activities and (ii) exceed specific capacity thresholds. These process-based definitions target, *inter alia*, pulp and paper mills, coke ovens, petroleum refineries, non-metallic mineral products (including the manufacture of glass, ceramics, and cement), and the manufacture of basic metals.<sup>3</sup> Indirect emissions are not taken into account, nor are electricity imports.

As is explained in Section 3.2 below, we match ETS installations to the manufacturing establishments operating them. An establishment identified in this way is considered as treated and referred to as an ETS plant. In the same vein, a firm is considered as treated and referred to as an ETS firm if it owns at least one ETS plant. As [Calel and Dechezleprêtre \(2016\)](#) and others have noted, the imperfect overlap in the definition of an installation used in the Emissions Trading Directive and the definition of plants and firms in official statistics induces variation in treatment status even among firms of similar size. This variation may arise because a fixed total capacity of a combustion or process-regulated activity at the plant level can be spread over multiple installations of varying capacity. As depicted in Figure ??

<sup>2</sup>To harmonize criteria across countries, as well as to include additional sectors, the directive was later amended (Directive 2009/29/EC)

<sup>3</sup>Beginning in 2012, emissions from other industries, such as aviation, have been included in the ETS as well.

example, a total of 30MW rated thermal input can be achieved with a single installation that, due to its size exceeding 20MW is regulated, or with a set of smaller installations of 15MW each that are not regulated.

In our baseline specification we shall consider treatment status at the firm level, as this internalizes within-firm spillovers in multi-plant firms and precludes double-counting. In further analysis, we shed light on those important aspects in order to obtain a more complete picture of the firms' behavioral response to incomplete regulation.

## 2.2 Related Literature

A growing empirical literature focuses on evaluating the effectiveness and costs of the ETS. A recent article by [Martin et al. \(2016\)](#) reviews this literature, in particular the early evaluation studies. Of particular relevance to our paper is a recent strand of microeconomic studies combining matching with a differences-in-differences estimator for evaluating the impact of the ETS on a variety of outcomes at treated firms. For instance, [Calel and Dechezleprêtre \(2016\)](#) analyze patent applications and find a small but significant positive impact on innovation among treated firms. [Petrick and Wagner \(2014\)](#), [Jaraite and Maria \(2016\)](#) and [Klemetsen et al. \(2016\)](#) study the impact of the ETS on CO<sub>2</sub> emissions using administrative data for manufacturing firms in Germany, Lithuania, and Norway, respectively. They find empirical support for emissions reductions among treated firms during the second phase of the ETS in some countries. In addition, none of these studies finds a detrimental effect of the ETS on the economic performance of participating firms. In a European-wide study based on balance-sheet data from multiple countries, [Marin et al. \(2017\)](#) provide evidence to suggest that the EU ETS did not negatively impact various measures of economic performance, including total factor productivity.<sup>4</sup> Using administrative data on German firms, [Lutz \(2016\)](#) also finds no significant negative effect of the ETS on revenue-based measures of productivity. Finally, [Borghesi et al. \(2016\)](#) test whether the ETS drove outward Foreign Direct Investments (FDI) in a large sample of Italian firms. They find increases in outward FDI on the intensive margin but no effect on the extensive margin.

Our paper contributes to this literature in a number of ways. First, we contribute to the strand of single-country studies of the ETS by providing new evidence on its environmental and economic performance impacts on French manufacturing firms. The French administrative data are rich in detail and of high quality. Furthermore, France is an interesting country to study because of its size and due to the fact that energy prices have been low compared to similar countries. Second, we exploit the fact that the French statistical authorities

---

<sup>4</sup>[Chan et al. \(2013\)](#) use similar data in a differences-in-differences analysis for selected energy-intensive sectors, but they do not match on observables to mitigate potential bias.

use a unique establishment identifier to perform high-quality entity linking across multiple datasets. This enables us to go beyond the existing research on various fronts. For instance, ours is the first paper to take the analysis to the sub-firm level, shedding light on treatment heterogeneities across plants and intra-firm leakage. Third, we link the core dataset to rich customs data on firm-level trade flows in order to assess possible leakage effects between firms in and outside the EU. By shedding light on various channels of leakage, this study makes an important step towards assessing the performance of market based instruments in the presence of incomplete environmental regulation.

## 3 Data

We compile a dataset of French manufacturing firms for the years from 1996 until 2012. This time interval covers several years prior to the announcement of the ETS in 2001, the announcement phase between 2001 and 2004, and and Phase I (2005-2007) and Phase II (2008-2012) after the implementation of the scheme. The data are obtained from various sources.<sup>5</sup>

### 3.1 Data sources

#### 3.1.1 Energy and emissions data

We obtain detailed fuel use data from the Annual Survey of Industrial Energy Consumption (EACEI),<sup>6</sup> a survey conducted annually by the French National Institute of Statistics and Economic Studies (INSEE - Institut National de la Statistique et des études économiques). The survey provides quantities and values of energy consumed by fuel type<sup>7</sup> - broadly speaking, electricity, steam, fossil fuels and biofuels - as well as by usages for each type of fuel.<sup>8</sup> Other variables available in the survey include geographical location and sectoral classification.

---

<sup>5</sup>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*.

<sup>6</sup>In French: *Enquête annuelle sur les consommations d'énergie dans l'industrie*.

<sup>7</sup>Information for the following fuel types is requested from the surveyed firms: electricity (bought, auto-produced and resold), vapor, natural gas, other types of gas available on the network, coal, lignite, coke, butane, propane, heavy fuel oil, heating oil, other petroleum products, the black liquor (a byproduct of the chemical decomposition of wood for making paper pulp), wood and its by-products, special renewable fuels, special non-renewable fuels.

<sup>8</sup>Electricity usages include: driving force, thermal uses, other uses (including electrolysis). For other types of energy, the survey distinguishes between manufacturing use, electricity production, raw materials use, heating use and other purposes.

Having reliable data on CO<sub>2</sub> emissions is of central importance to our study. We calculate emissions at both treated and untreated firms using the detailed energy consumption data from EACEI in conjunction with standardized conversion factors provided by the French Environment & Energy Management Agency (ADEME).<sup>9</sup> Consequently, a firm will only be in our core dataset if it reports detailed energy consumption data under the EACEI. The sampling frame for the EACEI includes all French manufacturing establishments,<sup>10</sup> and the response rate is close to 90 percent. This speaks to the high representativeness of the dataset, but it is important to note that not all establishments are covered, and that sampling rules have changed over time.

Until 2007, firms included in the EACEI survey were in sectors 12 to 37<sup>11</sup> according to the NAF rev.1 classification, equivalent at the two-digit level to the NACE rev.1. In more recent years, different sampling weights were applied to draw about 12,000 establishments for the sample. Specifically, the sample includes (i) all industrial establishments with 20 employees or more in the most energy consuming sectors<sup>12</sup>; (ii) all establishments with more than ten employees in sector 20.11Z (manufacturing of industrial gases); (iii) all establishments with more than 250 employees on the 31st of December of that year; (iv) a random sample of establishments with employment between 20 and 249 employees in sectors that are not energy intensive.

In sum, while the subsequent analysis is not based on the universe of French manufacturing firms, it draws on a database designed to provide a representative picture especially of the most energy intensive firms in French manufacturing while living up to the high standards of data collection for official statistics in France.

### 3.1.2 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 called Unified Corporate Statistics System (SUSE). We use two datasets that are based on SUSE, namely the FICUS database which covers the years from 1994 to 2007, and its successor FARE which provides data for the years from 2008 until 2012. These datasets provide general information about the

---

<sup>9</sup>ETS participants in France are required to use the same conversion factors when reporting their emissions.

<sup>10</sup>The level of survey is the establishment rather than the enterprise given that energy consuming materials, electricity and gas meters and fuel tanks are held at that level.

<sup>11</sup>The following sectors are excluded: 15 - Manufacture of food products and beverages, 20.1A, 22.1 and 23

<sup>12</sup>23.32Z - Manufacture of bricks, tiles and construction products, in baked clay; 23.51Z Manufacture of cement; 23.52Z- Manufacture of lime and plaster

firm (identifier, industry classification, head office address, total number of workers employed, age, etc.), the income statement (containing variables such as total turnover, total labor costs and value added) as well as balance sheet information (e.g. various measures of capital, debt and assets).<sup>13</sup> To measure capital, we used the value of the gross fixed tangible assets, which includes machinery, equipment and buildings.

### 3.1.3 Trade data

Trade data for the period of 1995 to 2012 are obtained from French Customs (DGDDI) and are reported at the firm level. The raw data are based on customs declarations that firms are required to submit, and which provide comprehensive annual records of value and quantity of exports/imports by destination/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) and includes the universe of trade flows from and to French firms, although reporting thresholds exist for compulsory declarations inside and outside the European Union. Outside the EU, exports or imports are only reported if their annual total is above €1,000 or 1,000 kg. Within the EU, these thresholds vary through time and by direction. In order to harmonize these different thresholds, we consider as non traders firms whose total exports or imports within or outside the EU are less than €150,000.

### 3.1.4 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.<sup>14</sup> Each 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 up 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 installations, which is calculated as the net balance of surrendered EUAs and verified emissions. This information is provided at the annual level.

---

<sup>13</sup>Observations displaying extreme growth in employment and value-added and emissions and emissions intensity are dropped.

<sup>14</sup>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 (CITL), operated by the Commission, which records and checks every transaction. Since 2012 the ETS registry has been operated in a centralized fashion as the EUTL.

## 3.2 Linking

The quality of the link between entities across data sets is an important determinant of the overall data quality in the empirical analysis to follow. Linking the EACEI, FICUS/FARE and trade data is straightforward as all three datasets use the SIREN (*Système d’Identification du Répertoire des Entreprises*) number as their identifier. Although each plant in the EACEI is identified by a SIRET number, the SIREN number corresponds to the first nine digits of the SIRET number. While the business data set is maintained by INSEE, the national ETS registry is managed by Caisse des Dépôts. The latter institution provides a link between the permit identifier (GIDIC) from the national registry and the SIREN identifier from INSEE. We combine this information with postcode identifiers to identify plants. When multiple plants of an ETS firm share the same postcode, we resort to name matching and latitude-longitude information available in the CITL in order to establish the linking to the corresponding plant in EACEI. Less than twenty manufacturing plants include more than one installation. In this way, we obtain an almost perfect link between manufacturing firms and plants and installations in the EUTL. We drop firms that we do not observe them either before the policy was announced in 2000 or after it was introduced in 2005 as this would prevent us from running our analysis.

## 3.3 Descriptive statistics

The resulting dataset includes 5,867 plants owned by 2,871 firms. 236 plants within 151 firms are part of the EU ETS. We expect that there is considerable treatment heterogeneity across plants, depending on whether a plant belongs to (i) a multi-plant firm with ETS plants only (fully-regulated firm), (ii) multi-plant firm with both ETS and non-ETS plants (partly-regulated firm), or (iii) a single-plant ETS firm. We observe that close to 70% of ETS plants are part of firms with non-ETS facilities and represent close to 60% of total emissions by ETS firms in 2000. Within the subset of firms that are partly regulated, 65% of the plants are not regulated by the ETS.

The first column of Table 1 shows how ETS and non-ETS firms differ with respect to the main variables that are available from our different datasets. We can see that ETS firms are on average larger, regardless of whether size is measured in terms of employment, value added, capital or imports. They also emit more CO<sub>2</sub> emissions and are more carbon intensive, as can be seen in the significant and large coefficients for log CO<sub>2</sub> and for the coal share. The following section describes how we construct a representative control group to compare regulated and non-regulated firms and plants given the observed differences in covariates.

## 4 Research Design

In this study, we exploit variation in the selection criterion by which firms are required to join the ETS. Building on the potential outcomes framework commonly used in the program evaluation literature, we propose that firms can be in one of two states: either part of the market-based ETS, or prevailing in a state of business as usual.<sup>15</sup>

### 4.1 Semi-Parametric Differences-in-Differences

Let  $ETS_i = 1$ , if firm  $i$  is a member of the ETS and is therefore part of the treatment group. Firm  $i$  is part of the ETS if it owns at least one installation that is regulated under this policy. Let  $ETS_i = 0$  if firm  $i$  is not part of the ETS and is therefore part of the control group. The potential outcomes  $Y_{it}(1)$  and  $Y_{it}(0)$ , conditional on membership and non-membership respectively, denote the outcome variables of interest for installation  $i$  in the post-treatment period ( $t = 1$ ) or the pre-treatment period ( $t = 0$ ). We are interested in estimating the average treatment effect on the treated (ATT)

$$\alpha_{ATT} = \mathbb{E}[Y_{i1}(1) - Y_{i1}(0)|ETS_i = 1] \quad (1)$$

where  $\alpha_{ATT}$  measures the average effect of the ETS on the outcome variable of interest.

The problem in identifying the causal effect of the ETS arises from missing data. Firm-level emissions data for ETS participants, during the years following the implementation of the program, can be used to identify  $\mathbb{E}[Y_{i1}(1)|ETS_i = 1]$ . However,  $\mathbb{E}[Y_{i1}(0)|ETS_i = 1]$  is not observed. To address this issue, counterfactual outcomes are constructed using emissions observed at firms that are not subjected to the ETS for the duration of the study.

Constructing such counterfactual outcomes is one of the key challenges in empirical research with non-experimental data. A naive estimate of the  $\alpha_{ATT}$  is obtained by computing unconditional differences-in-differences for ETS firms and non-ETS firms. Even when controlling for observables, the differences-in-differences estimator may attribute some of the changes between the outcome variables to the ETS when they could really be the result of other systematic differences between treatment and control plants.

In order to reduce the bias induced by systematic differences between ETS and non-ETS firms, we combine a semi-parametric conditioning strategy with a differences-in-differences approach. This is the idea behind the generalized difference-in-difference estimator suggested

---

<sup>15</sup>See Holland (1986) for a deeper discussion of causal inference, the potential outcomes framework, and its history.

by Heckman et al. (1997, 1998):

$$\begin{aligned}\alpha_{ATT}^{matched} &= \mathbb{E}[Y_{i1}(1) - Y_{i0}(1)|X_i, ETS_i = 1] \\ &= \frac{1}{N_1} \sum_{j \in I_1} \left\{ (Y_{jt_1}(1) - Y_{jt_0}(0)) - \sum_{k \in I_0} \omega_{jk}(X_j, X_k) \cdot (Y_{kt_1}(0) - Y_{kt_0}(0)) \right\} \quad (2)\end{aligned}$$

where  $I_1$  denotes the set of ETS firms,  $I_0$  the set of non-ETS firms, and  $N_1$  the number of participating firms in the treatment group. The treated firms are indexed by  $j$ ; the control firms are indexed by  $k$ . The weight placed on a non-ETS firm when constructing the counterfactual estimate for ETS firm  $j$  is  $\omega_{jk}$ . These weights can be calculated using any matching approach.

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

$$(Y_{j,t} - Y_{j,2000}) - (Y_{k,t} - Y_{k,2000}) = \sum_{\tau=1}^4 \sum_{t=1996}^{2012} \beta_{\tau} \times \mathbb{1}\{t \in \Theta_{\tau}\} + \varepsilon_{jt}$$

where

$\Theta_1 =$	$\{1996, \dots, 1999\}$	Pre-announcement period
$\Theta_2 =$	$\{2001, \dots, 2004\}$	Announcement period
$\Theta_3 =$	$\{2005, 2006, 2007\}$	Trading Phase I
$\Theta_4 =$	$\{2008, \dots, 2012\}$	Trading Phase II.

The LHS of Equation (4.1) denotes the difference in outcome between treated firm  $j$  and matched control firm  $k$  in year  $t$  relative to the base year 2000, i.e. just before the announcement of the EU ETS. The coefficients of interest are  $\beta_{\tau} = \alpha_{ATT}^{matched}$  for Phase  $\tau$ , which provides the relative effect of the EU ETS on regulated firms in phase  $\tau$ , compared to the year 2000.

They capture the effect of the ETS on treated plants, relative to the matched control plant, in each year relative to the year 2000.

## 4.2 Implementation

This section explains how we construct counterfactual outcomes for treated firms in the ETS by matching them to control firms that are observationally similar. The discussion relates

to our main specification. The appendix provides alternative specifications and robustness tests.

#### 4.2.1 Matching Variables

We follow Fowlie et al. (2012) and match treated firms to control firms based on (i) the logarithm of CO<sub>2</sub> emissions in the year 2000 (the year prior to the announcement of the ETS) and (ii) on the exact 2-digit sector of the firm.

Some firms in our sample own plants that engage in different activities, as specified by their NACE sectoral classification. To account for this, we define a new sector variable SUPERNACE at the firm level which is based on the combination of these plant-level activities. For example, if a firm owns two plants and both produce in NACE 12, then the SUPERNACE is 12 and the firm would be matched to a control firm in the same sector (with SUPERNACE 12). In contrast, for a firm with one plant producing in NACE 12 and another one in NACE 17, we define SUPERNACE to be 1217 and match it to a control firm within SUPERNACE 1217 (where the ordering of sectoral codes does not matter, e.g., SUPERNACE “1217” is equivalent to SUPERNACE “1712”).

The set of matching variables is chosen to produce a comparison group that has similar characteristics as the treatment group while maximizing the number of successful matches.<sup>16</sup> This parsimonious matching strategy allows us to verify that covariates are balanced between treatment and control firms across both matched and unmatched variables. Moreover, rather than matching on pre-treatment trends we let the data speak to the validity of the assumption that pre-treatment trends in the outcome variables are parallel. The exact match on sector controls for sector specific shocks to the outcome variables post treatment.

#### 4.2.2 Sample Restrictions

Following the above matching steps, we impose further restrictions on the sample in order to improve the plausibility of the set of control firms. We impose a maximum nearest-neighbor distance by restricting matched control firms to be within the 95 percentile of the distance distribution. Moreover, we restrict the difference in the number of plants between treatment and control firms to be lower than four. For example, a match between a treatment firm with 10 plants to a control firm with 7 plants or 13 plants would be included in our analysis sample. However, a match between a treatment firm with 10 plants to a control firm with 6 plants or 14 plants would be excluded.

---

<sup>16</sup>Appendix A presents balance tests and results based upon alternative matching specifications.

Table 1 provides pre-match and post-match balance tests which highlight the improvement in common support that is achieved through the matching process and sample restrictions. Figures A.1 and A.2 in Appendix A present graphical evidence analogous to the balance tests. Our matching does not completely eliminate differences between treated and controls. Having panel data helps in this regard as it allows us to difference out persistent heterogeneity between treated and control firms.

### 4.2.3 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 (2016) demonstrate that bias in the estimation of the variance can occur if the covariates in the regression are correlated with the error term, conditional on the variables that have been matched on. They demonstrate that valid inference can be conducted if units one matches *without replacement* and clusters standard errors at the level of the match.

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

Given the trade-off between minimizing bias in the point estimates vs. in the standard errors, we give priority to the former and use matching with replacement in our main specification. To address the point made by Abadie and Spiess (2016), we cluster in two ways. The first cluster is at the level of the match (the firm) and also addresses serial correlation. The second cluster is at the control-firm-year level to account for correlation across observations that are matched to the same control observation. This second cluster is redundant in the event that each treatment firm is matched to a unique control firm.

### 4.2.4 Empirical support for identifying assumptions

Identifying the average treatment effect on the treated with the research design presented above is predicated on a number of assumptions, some of which are straightforward to examine in the data. For instance, the common support condition can be evaluated by studying

Table 1: Difference-in-Means between Regulated and Unregulated Firms

	(1)	(2)
	PRE-MATCH DIFFERENCE	POST-MATCH DIFFERENCE
	(Full Sample)	(Matched Sample)
log (CO <sub>2</sub> )	2.284*** (0.126)	0.359** (0.175)
log (EMPLOYMENT)	0.620*** (0.096)	0.325 (0.269)
log (VALUE ADDED)	0.081*** (0.106)	0.265 (0.218)
log (CAPITAL)	1.175*** (0.129)	0.406* (0.210)
log (IMPORTS)	0.961*** (0.164)	0.578* (0.316)
log (IMPORTS) (Non-EU)	1.364*** (0.195)	0.824* (0.442)
log (IMPORTS) (EU)	0.851*** (0.166)	0.337 (0.319)
(COAL SHARE)	0.0206* (0.0111)	0.00741 (0.0135)
(GAS SHARE)	0.080*** (0.020)	-0.063 (0.081)
(ELECTRICITY BOUGHT SHARE)	-0.034*** (0.010)	-0.025 (0.248)

NOTES: Significance levels are indicated as \* 0.10 \*\* 0.05 \*\*\* 0.01. Presented coefficients report the difference in outcome variables between treatment and control firms. Column 1 presents the average difference between unmatched treatment and control firms. Column 2 presents the average difference between matched treatment and control firms. Distance restrictions between treatment and control plants are imposed at the 95th percentile. Standard errors are clustered at the firm-level (Pre-Match) and two-way clustered by firm and matching group (Post-Match).

density plots of the kind presented in Appendix A.1 or even enforced by the matching algorithm.

What is more, the assumption that outcomes across groups follow parallel trends in the absence of treatment can be falsified by looking at data. In the results to follow we report a placebo treatment effect during the pre-announcement period, which allows the reader to assess the plausibility of that assumption one specification at a time.

Another assumption commonly made in program evaluation is that the potential outcomes at one firm are independent of the treatment status of other firms – the stable unit treatment value assumption (SUTVA; Rubin, 1980). SUTVA is necessary to rule out spillovers and general equilibrium effects. Although SUTVA is not testable, researchers have successfully evaluated its plausibility through indirect tests and provided some insight into the sign of the bias (Fowlie et al., 2012). Spillovers are potentially important in our application, as it might be the case that within a firm, emissions are reallocated from ETS plants to non-ETS plants, or that the policy may result in general equilibrium effects between treatment and control firms. If the ETS affects the productivity of larger firms, non-regulated firms may increase their market share, increasing output and emissions. By aggregating our analysis to the firm level we are able to internalize any spillovers within-firm; however, we are unable to rule out spillovers between firms *a priori*.

## 5 Results

### 5.1 Firm-level Results

We begin by estimating the effects of the ETS on firm-level emissions, inputs (employment and tangible assets) and value added, using the research design discussed previously. The results of this exercise are presented in Table 2 and Figure 2.

The first column of Table 2 reports the treatment impact on CO<sub>2</sub> emissions, the outcome targeted by the ETS, for different time periods. We cannot reject the hypothesis that there were no differences in emissions trends across groups prior to the announcement of the ETS. However, following its announcement, we see that *ex-post* regulated firms increased their CO<sub>2</sub> emissions by 6.6 percentage points relative to *ex-post* unregulated firms. This increase could be motivated by rent seeking. If regulated firms correctly anticipated that they were going to receive free permits in proportion to historical emissions under a grandfathering allocation scheme, they could extract large permit rents by *ramping up* emissions during the

Table 2: The Effect of the EU ETS on Firm Outcomes

	(1)	(2)	(3)	(4)
	$\Delta \log(\text{CO}_2)$	$\Delta \log(\text{VA})$	$\Delta \log(\text{Emp})$	$\Delta \log(\text{Capital})$
PRE-ANNOUNCEMENT	0.011 (0.020)	0.033 (0.027)	0.022 (0.019)	0.017 (0.027)
ANNOUNCEMENT PERIOD	0.066*** (0.025)	0.050 (0.038)	0.011 (0.025)	0.0046 (0.028)
TRADING PHASE I	0.001 (0.045)	-0.023 (0.051)	-0.045 (0.038)	-0.033 (0.046)
TRADING PHASE II	-0.135** (0.059)	-0.005 (0.098)	-0.075 (0.056)	0.077 (0.083)
Mean of exp(Dep. Var.) in 2000	85.83	53,515	668	132,089
Observations	1,575	1,575	1,575	1,542

NOTES: Significance levels are indicated as \* 0.10 \*\* 0.05 \*\*\* 0.01. Distance restrictions between treatment and control firms are imposed at the 95th percentile. Standard errors are two-way clustered by firm and matching group.

announcement phase.<sup>17 18</sup>

During Trading Phase I of the ETS, we find that there are no differences in emissions between treated and control firms, suggesting that Phase I was ineffective in reducing emissions by more than what would have happened in absence of treatment. This is consistent with the fact that the system on aggregate was over-allocated and thus compliance was not costly (Ellerman and Buchner, 2008). Furthermore, the point estimate implies that the ramping-up effect from the announcement phase has been neutralized.

Moving on to Phase II, we find that regulated firms reduced their emissions by 13.5%, relative to the year 2000, compared to unregulated firms. This result suggests that the ETS has been rather effective in its prime objective to reduce emissions and calls for further investigation of the pathways of emissions abatement.

A look to economic performance outcomes, reported in the remaining columns of Table 2, provides a first insight. The estimated coefficients for value added and employment, reported in columns two and three, respectively, are all statistically insignificant at conventional levels. During the pre-treatment phase, we cannot reject the hypothesis of equal trends in valued added and employment across treated and control firms. The ramping-up

<sup>17</sup>Under the French National Allocation Plan for trading phase I, free permits were granted to industrial installations in proportion emissions during the period from 1996 until 2002. For new installations, data from 2004 and 2005 would also be taken into account. Cf. [Ministère de l'environnement \(2005\)](#)

<sup>18</sup>Bushnell et al. (2013) present evidence that free permit endowments were highly valued by investors of publicly traded firms that were regulated under the ETS.

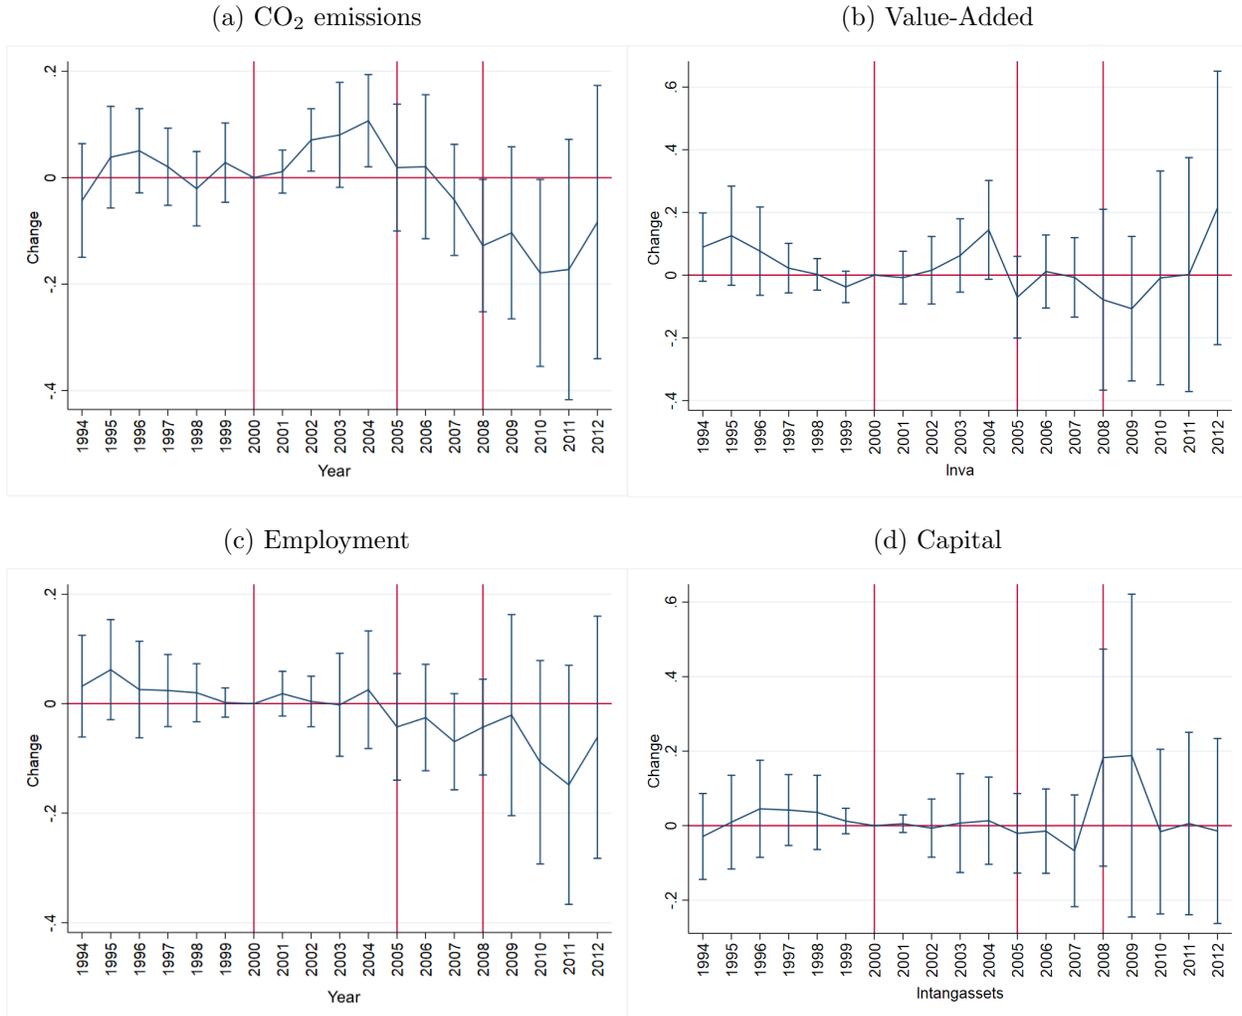
of emissions observed during the announcement period is mimicked by a point estimate for value added of similar magnitude; however, this is statistically insignificant. After that, the changes in value added are disentangled from changes in CO<sub>2</sub> emissions, suggesting that emissions abatement is not simply brought about by a reduction in the scale of production. The point estimates for employment are negative in both trading phases. However, the large standard errors imply that the null hypothesis of no employment impact cannot be rejected.

The last column of Table 2 reports the effect of the EU ETS on tangible assets. Investment in capital is a plausible mechanism for actual emissions abatement, because regulated firms might invest in more efficient machinery and boilers to bring down carbon emissions while maintaining output levels constant. The point estimates we obtain fluctuate around zero and are quite noisy (reflecting the lumpiness in adjustments to capital). Again we cannot reject the Null hypothesis that there is no significant difference between the accumulation of tangible assets at treated and control firms. This suggests that either the abatement investments made were too small to be salient in the overall investments made by firms, or that firms engaged in other margins of adjustment to reduce emissions.

Figure 2 reports annual treatment effects on CO<sub>2</sub> emissions, employment, value added, and capital, along with 95%-confidence bands. The increase in emissions during the announcement period as well as the decrease in Phase II are apparent. The annual coefficients are estimated with more noise than the period averages reported in Table 2. For value added, employment and capital, none of the annual treatment effects is estimated with precision. Our results also provide first indications regarding concerns about economic performance and carbon leakage. Firstly, if a regulated firm becomes less competitive in the product market we would expect to see a reduction in output. Secondly, if a regulated firm abates carbon emissions by out-sourcing carbon intensive parts of the value chain we would expect an increase in intermediate inputs. Both adjustments should reduce value added, but we find no such effect.

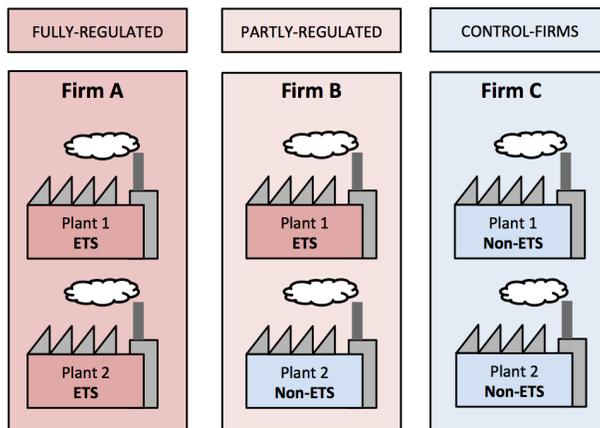
While this would suggest that the emission reductions we find translate into abatement in the global sense, we cannot rule out that there is no other reallocation of emissions within France, either between regulated and non-regulated installations within-firm, or between regulated and unregulated firms. We will examine some of the possible channels and margins of adjustment more directly in the following sections, firstly by examining partially regulated multi-plant firms - i.e. firms with both ETS and non-ETS installations - within France and secondly by looking at imports of intermediates from abroad.

Figure 2: Changes in outcome variable at treated vs. control firms



Notes: The four sub-figures represent the differences in  $\log(\text{CO}_2 \text{ emissions})$ ,  $\log(\text{Value added})$ ,  $\log(\text{Employment})$ , and  $\log(\text{Capital})$  between treatment and control firms over time, with the same matching approach as for the results in Table 2 and with 2000 as the base year. Confidence interval bars at 95%. The red lines divide the sample into four phases: (i) pre-announcement, (ii) announcement period, (iii) Trading Phase I and (iv) Trading Phase II (from left to right).

Figure 3: Regulation status at firm level



Source: Own.

## 5.2 Unbundling Firm-level Emissions Reductions

In this section we explore the degree to which emission reductions may have arisen as a result of firm-level adjustments to the organization of the production process that shift emissions elsewhere instead of avoiding them. Specifically, we hypothesize and empirically test two such adjustments. First, firms may reallocate emissions from regulated to unregulated facilities within the firm. This behavior would imply that the firm-level estimate of abatement reported in Table 2 is smaller than what it would have been in the absence of intra-firm emissions shifting. Second, firms may adjust their supply chains, resulting in a reallocation of “dirty” production from regulated firms to unregulated firms. While we are unable to look at domestic sourcing of intermediate inputs, we shall empirically examine the extent to which this effect operates through international trade.

### 5.2.1 Fully regulated ETS firms vs. partly regulated ETS firms

We begin by exploring heterogeneity in the treatment effects of the EU ETS with respect to the treatment intensity, which determines the scope for within-firm reallocation. To this end, we define fully regulated firms as firms that own only ETS plants, and partly regulated firms as firms that own at least one unregulated plant and thus have the ability to remove emissions from the ETS through within-firm reallocation (see Figure 3 for an illustration). We then re-estimate Equation 4.1 and interact the treatment effect for each period with a dummy for fully regulated firms.

The results are presented in Table 3 which reports the treatment effects for each group.<sup>19</sup>

<sup>19</sup>For better comparability, the table reports total effects for both groups rather than level effects for

Table 3: Heterogenous Effects of the EU ETS on Firm Outcomes

	(1)	(2)	(3)	(4)
	$\Delta \log(\text{CO}_2)$	$\Delta \log(\text{VA})$	$\Delta \log(\text{Emp})$	$\Delta \log(\text{Capital})$
<i>A. Fully regulated firms</i>				
PRE-ANNOUNCEMENT	0.001 (0.025)	0.071** (0.032)	0.022 (0.022)	0.084 (0.054)
ANNOUNCEMENT PERIOD	0.044 (0.032)	0.039 (0.045)	0.006 (0.025)	-0.053 (0.038)
TRADING PHASE I	-0.077* (0.044)	-0.090 (0.060)	-0.051 (0.039)	-0.062 (0.084)
TRADING PHASE II	-0.231*** (0.062)	-0.050 (0.110)	-0.088 (0.062)	0.031 (0.134)
Mean of exp(Dep. Var.) in 2000	92,128	32,868	439	91,237
<i>B. Partly regulated firms</i>				
PRE-ANNOUNCEMENT	0.010 (0.029)	-0.028 (0.043)	0.021 (0.032)	-0.034 (0.049)
ANNOUNCEMENT PERIOD	0.097*** (0.033)	0.068 (0.053)	0.020 (0.040)	0.036 (0.039)
TRADING PHASE I	0.114 (0.072)	0.077 (0.078)	-0.036 (0.067)	0.004 (0.077)
TRADING PHASE II	0.0047 (0.090)	0.060 (0.139)	-0.057 (0.088)	0.058 (0.086)
Mean of exp(Dep. Var.) in 2000	76,256	84,899	1,016	194,186
Observations	1,575	1,575	1,575	1,542

NOTES: Significance levels are indicated as \* 0.10 \*\* 0.05 \*\*\* 0.01. Distance restrictions between treatment and control firms are imposed at the 95th percentile. Coefficients reported in panels A and B are obtained in a single regression for each column, with treatment interactions for fully regulated. These coefficients are displayed as the level effect for a given period plus the interaction term, e.g. Phase II + Phase II  $\times$  (Fully Regulated). Standard errors are two-way clustered by firm and matching group.

The first column reveals a fundamental difference in how firms with different regulation status respond to the ETS. While the group of fully regulated firms drives the emissions reductions found in Table 2 above, the group of partly regulated firms drives the ramping-up effect also documented there. Specifically, panel A shows that fully regulated firms reduced their CO<sub>2</sub> emissions during both trading phases, notably in Phase II by 23% relative to unregulated firms. As was found already for the full sample, this abatement is accompanied by no significant change in value added (column 2), employment (column 3), or capital (column 4).<sup>20</sup>

In contrast to this, panel B shows that partly regulated firms ramped up emissions by 9.7% during the announcement period, presumably in the hopes of obtaining higher allocations of free permits at a later trading stage. No significant reduction occurs during the trading phases as these firms have the ability to reallocate emissions and economic activity across plants within firm. The heterogeneity in treatment effects across firms that have different opportunities for shifting emissions within-firm in ways that remove them from the cap highlights that this kind of leakage could be potentially important. We shall examine this hypothesis in the next section in the framework of a plant-level analysis of the impact of the EU ETS on emissions.

Before doing so, we explore the heterogeneity of the treatment effects across firms of different treatment intensity with respect to abatement via fuel switching. Table (4) reports coefficient estimates from Equation 4.1 modified to include an interaction term for fully regulated firms, and for four relevant outcomes. The first three columns focus on the overall fuel mix by reporting the change in the share of coal, natural gas, and electricity, respectively, in total energy consumption. The dependent variable in column 4 is (the change in) the share of electricity bought from the grid in total electricity consumption. For fully regulated firms (panel A) we do not find any significant changes to their fuel mix in any of the periods. However, these firms increase the share of bought electricity by 6.2 percentage point in Trading Phase II. That is, while holding overall electricity consumption constant (or decreasing it if the imprecise point estimate in column 3 is taken at face value), this type of firm switches from own generation to grid electricity, 85% of which is generated from zero-carbon sources in France. In contrast, we find no such effect for partly regulated firms. This was to be expected having found that the average firm in this group does not abate CO<sub>2</sub> emissions. The only significant coefficient arises during the announcement period, where these firms increase the use of electricity in their fuel mix by 5.4 percentage point. To the extent that this

---

partly regulated firms and interaction effects for fully regulated firms.

<sup>20</sup>Fully regulated firms do exhibit stronger increases in value added during the pre-treatment period, however.

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

	(1)	(2)	(3)	(4)
	$\Delta(\text{Fuel shares in total energy use})$			$\Delta\left(\frac{\text{Electricity bought}}{\text{Total electricity}}\right)$
	Coal	Natural gas	Electricity	
<i>A. Fully regulated firms</i>				
PRE-ANNOUNCEMENT	0.002 (0.006)	0.029 (0.025)	-0.002 (0.007)	0.012 (0.023)
ANNOUNCEMENT PERIOD	0.007 (0.032)	0.023 (0.026)	-0.036 (0.025)	0.019 (0.015)
TRADING PHASE I	0.002 (0.01)	0.044 (0.039)	-0.037 (0.025)	0.006 (0.014)
TRADING PHASE II	0.006 (0.014)	0.073 (0.051)	-0.046 (0.032)	0.062** (0.025)
<i>B. Partly regulated firms</i>				
PRE-ANNOUNCEMENT	0.0002 (0.0052)	-0.0183 (0.0229)	0.0015 (0.0065)	-0.0263 (0.0221)
ANNOUNCEMENT PERIOD	-0.0002 (0.0328)	0.0093 (0.0208)	0.0540** (0.0227)	0.0038 (0.0060)
TRADING PHASE I	-0.0152* (0.0091)	0.0100 (0.0327)	0.0221 (0.0235)	0.0107 (0.0073)
TRADING PHASE II	-0.0109 (0.0112)	-0.0604 (0.0419)	0.0274 (0.0309)	-0.0113 (0.00916)
Observations	1,575	1,575	1,575	1,573

NOTES: Significance levels are indicated as \* 0.10 \*\* 0.05 \*\*\* 0.01. Distance restrictions between treatment and control firms are imposed at the 95th percentile. Coefficients for fully regulated firms are displayed as the level effect + the interaction term, e.g. Phase II + Phase II  $\times$  Fully Regulated. Standard errors are two-way clustered by firm and matching group.

electricity is generated on site with dirty energy sources, this can account for the ramping-up of CO<sub>2</sub> emissions found for this group of firms.<sup>21</sup>

### 5.2.2 Plant-level Analysis of Within-Firm Reallocation

We have adopted a firm-level approach to studying the impacts of the ETS not only because it is instrumental in the construction of a credible counterfactual via matching, but also because it accounts for centralized decision-making across plants that belong to the same firm and hence avoids double counting of emissions reductions across those plants. However, when

<sup>21</sup>In Table B2 in Appendix B we report the corresponding coefficients estimated without the interaction. While the ramping-up effect is present therein as well, the impact on the electricity bought share is masked by the underlying heterogeneity in this regression.

it comes to partly regulated firms, the firm-level regressions mask important heterogeneities in the way that plants with different regulatory status respond to the ETS. To shed light on this, we re-estimate Equation (4.1) at the plant level. In keeping with the view that decision-making is not independent across plants that belong to the same firm, the matching of treatment and control plants continues to be implemented at the firm level. That is, we match every treated plant with the most suitable control plant from the same firm that we used in the results above. This also maintains close comparability with the results obtained above.

Table 5 reports the results obtained by this exercise. When including all 1,625 plants in the regression (column 1) we find evidence of a (weakly significant) increase in emissions during the announcement period, as well as significant reduction in emissions by 19% during Phase II. These results are qualitatively similar to the firm-level results, except that the ramping up effect is only marginally significant.

In column 2 we restrict our attention to ETS plants that belong to fully regulated firms, including single-plant ETS firms. These plants are part of firms that have little scope for within-firm leakage. Consistent with this premise, we estimate significant reductions in Phase I (-12.7%) and Phase II (-24.1%), very similar to the corresponding firm-level estimates reported in Panel A of Table 3.

The last three columns of Table 5 report treatment effects for plants in firms with both regulated and unregulated plants. These plant-level estimates are thus underlying the firm-level estimates reported in panel B of Table 3. The main finding is that the ramping up effect during the announcement period is statistically significant at regulated plants but not at unregulated plants. This is consistent with the rent-seeking motivation we have conjectured earlier, because emissions at non-ETS plants have no effect on future permit allocations. The results for the trading phases are less clear cut. We do not estimate a statistically significant effect on emissions for either type of plant. This is consistent with the results obtained in panel B of Table 3, where we estimated insignificant precise zero reductions in emissions during Phase II. Moreover, the fact that the point estimates are of opposite sign in column (4) is consistent with intra-firm reallocation of emissions. For instance, in Trading Phase II the point estimate for regulated plants in partly regulated firms implies less than half as much abatement than in fully regulated plants, while the one for unregulated plants implies an increase in emissions by 15.1%. While these plant-level estimates suffer from a lack of precision due to the limited sample size, they support the view that the installation-based regulatory approach adopted under the ETS gives rise to a fundamental treatment heterogeneity across plants of the same firm. In column (5) we assign unregulated plants in regulated firms treatment status, alongside their already treated plants. We find that this

pooled plant treatment results in no significant reduction in emissions, consistent with the findings in panel B of Table 3.

Table 5: Within-Firm Reallocation? Plant-level Evidence

	(1)	(2)	(3)	(4)	(5)
	All	Fully regulated firms	Plants in partly regulated firms	ETS	Non-ETS
			ETS	Non-ETS	Both
	Dependent variable is $\Delta \log(\text{CO}_2)$				
PRE-ANNOUNCEMENT	0.055 (0.035)	0.071 (0.046)	0.029 (0.052)	0.101 (0.068)	0.053 (0.044)
ANNOUNCEMENT PERIOD	0.050* (0.026)	0.038 (0.038)	0.068* (0.036)	0.057 (0.078)	0.060 (0.038)
TRADING PHASE I	-0.053 (0.046)	-0.127*** (0.037)	0.063 (0.093)	-0.103 (0.148)	-0.006 (0.071)
TRADING PHASE II	-0.190** (0.0924)	-0.241** (0.0945)	-0.108 (0.150)	0.151 (0.182)	-0.069 (0.127)
Mean of CO <sub>2</sub> in 2000	65.79	85.322	35.135	20.061	33.011
Observations	1,625	999	628	469	1,086

NOTES: Significance levels are indicated as \* 0.10 \*\* 0.05 \*\*\* 0.01. Distance restrictions between treatment and control plants are imposed at the 95th percentile. Standard errors are two-way clustered by firm and matching group.

In light of these findings, it seems likely that overall emission reductions would have been substantially higher in the absence of within-firm reallocation. As partly regulated firms represent close to 60% of regulated emissions, their ability to reallocate emissions to unregulated plants undermines the overall effectiveness of the ETS. The question that remains is whether the emission reductions within fully regulated firms are real, or the result of between-firm leakage.

### 5.3 Emissions and International Trade

A specific channel of between-firm reallocations that we are able to empirically analyze is the extent to which the ETS induced increases in imports from outside of the ETS. We test this hypothesis by fitting Equation (4.1) to firm-level import data. Table 6 presents the results of this exercise. In column 1 we estimate the impact of the ETS on total imports, finding no significant effects. Columns two and three break up import flows according to whether they come from countries that participate in the ETS and from those that are not part of the ETS. In either case, we cannot reject the Null hypothesis that the ETS had no

effect on imports. On balance, this evidence points to a rather limited role that between-firm reallocations through trade may have played in explaining the estimated emission reductions. These findings hold whether we look at all firms or restrict our attention to fully regulated ETS firms.

Table 6: Carbon Trading or Trading Carbon? Firm-Level Imports

	(1)	(2)	(3)
	Total	ETS Countries	Non-ETS Countries
	$\Delta \log(1+\text{Imports})$	$\Delta \log(1+\text{Imports})$	$\Delta \log(1+\text{Imports})$
PRE-ANNOUNCEMENT	-0.124 (0.098)	-0.224 (0.168)	-0.169 (0.124)
ANNOUNCEMENT PERIOD	0.0048 (0.127)	-0.0358 (0.129)	-0.095 (0.200)
PHASE I	-0.014 (0.158)	-0.142 (0.165)	0.254 (0.242)
PHASE II	-0.199 (0.197)	-0.111 (0.210)	-0.585 (0.486)
Mean of exp(Dep. Var.) in 2000	49,563	23,594	25,969
Observations	1,610	1,610	1,610

NOTES: Significance levels are indicated as \* 0.10 \*\* 0.05 \*\*\* 0.01. Distance restrictions between treatment and control firms are imposed at the 95th percentile. Standard errors are two-way clustered by firm and matching group.

Furthermore, we test whether the ETS changed imports at the extensive margin by analyzing their portfolio of imports. The results are reported in Table 7 and exhibit no changes in the number of products, or countries that the firm imports from. Thus, if there are any adjustments to imports, they must be happening within the same portfolio, along the intensive margin.

Collectively, this evidence suggests that there is limited between-firm reallocations of emissions occurring through international trade. This is encouraging as it suggests that the estimated emission reductions in our main results contribute to emissions abatement at the global scale. However it may also be the case that between-firm reallocations are limited because firms engage in within-firm reallocations at a lower cost. As discussed, the overall reduction in emissions would have been substantially larger in the absence of within-firm reallocations, assuming that firms would not then engage in between-firm reallocations as a substitutable avoidance strategy.

Table 7: Carbon Trading or Trading Carbon? Changes to the Composition of Firm-Level Imports

	(1) Total	(2) ETS Countries	(3) Non-ETS Countries
A. Dep. Var. is $\Delta\#$ of trading partners			
PRE-ANNOUNCEMENT	-0.424 (0.349)	-0.265 (0.170)	-0.159 (0.280)
ANNOUNCEMENT PERIOD	-0.367 (0.453)	-0.172 (0.233)	-0.195 (0.340)
TRADING PHASE I	-0.364 (0.548)	-0.407 (0.317)	-0.043 (0.496)
TRADING PHASE II	-0.797 (0.869)	-0.703 (0.431)	-0.094 (0.593)
Mean of Dep. Var. in 2000	14.74	8.23	6.51
Observations	1,610	1,610	1,610
B. Dep. Var. is $\Delta\#$ of products			
PRE-ANNOUNCEMENT	-0.365 (2.418)	-3.206 (3.205)	-0.317 (1.646)
ANNOUNCEMENT PERIOD	1.597 (2.870)	3.577 (2.280)	-2.351 (1.831)
TRADING PHASE I	4.168 (4.048)	3.896 (3.303)	-0.268 (2.380)
TRADING PHASE II	7.800 (5.115)	5.610 (4.404)	2.842 (2.537)
Mean of Dep. Var. in 2000	75.23	53.77	33.84
Observations	1,610	1,610	1,610

Notes: Significance levels are indicated as \* 0.10 \*\* 0.05 \*\*\* 0.01. Distance restrictions between treatment and control firms are imposed at the 95th percentile. Standard errors are two-way clustered by firm and matching group.

## 6 Conclusion

In this paper quasi-experimental evidence on the environmental consequences of the European Union Emissions Trading Scheme is presented, with a focus on understanding the degree to which firm behavior can undermine effectiveness.

Using rich administrative data on French manufacturing firms we first match regulated firms to unregulated firms, using pre-policy announcement characteristics. For the matched sample we subsequently adopt a difference-in-difference identification strategy to evaluate the causal effects of the EU ETS on regulated firms.

We find that regulated firms reduced their emissions by 13.5% during the second phase of the EU ETS, compared to unregulated firms. In addition, we find no statistically significant effects of the EU ETS on firms' value added, employment, or capital. In contrast, abatement at fully regulated firms is achieved through switching into low-carbon electricity from the grid. Furthermore, we provide evidence to suggest that at least part of our estimated emission reductions are likely to be real, in a global sense, as regulated firms do not make changes to their supply chains through international trade. In addition, the absence of significant changes to value-added is consistent with the view that the same is likely to be true for intermediate inputs from domestic sources.

Notwithstanding this, firms also appear to make use of loopholes to circumvent the regulation, especially by reducing emissions in regulated plants and increasing emissions in unregulated plants. Consequently, aggregate emission reductions could have been substantially larger in the absence of within-firm reallocation. In sum, our findings suggest that emissions trading can be an effective tool for reducing greenhouse gas emissions, but that opportunistic behavior by firms can undermine its efficacy.

## References

- ABADIE, A. AND J. SPIESS (2016): "Robust Post-Matching Inference," Mimeograph, Harvard University.
- BORGHESI, S., C. FRANCO, AND G. MARIN (2016): "Outward Foreign Direct Investments Patterns of Italian Firms in the EU ETS," SEEDS Working Papers 0116, SEEDS, Sustainability Environmental Economics and Dynamics Studies.
- BUSHNELL, J. B., H. CHONG, AND E. T. MANSUR (2013): "Profiting from Regulation: Evidence from the European Carbon Market," *American Economic Journal: Economic Policy*, 5, 78–106.

- CALEL, R. AND A. DECHEZLEPRÊTRE (2016): “Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market,” *The Review of Economics and Statistics*, 98, 173–191.
- CHAN, H. S. R., S. LI, AND F. ZHANG (2013): “Firm competitiveness and the European Union Emissions Trading Scheme,” *Energy Policy*, 63, 1056 – 1064.
- EATON, J., S. KORTUM, AND F. KRAMARZ (2011): “An Anatomy of International Trade: Evidence From French Firms,” *Econometrica*, 79, 1453–1498.
- ELLERMAN, A. D. AND B. K. BUCHNER (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, D., C. MARCANTONINI, AND A. ZAKLAN (2016): “The European Union Emissions Trading System: Ten Years and Counting,” *Review of Environmental Economics and Policy*, 10, 89–107.
- FOWLIE, M., S. P. HOLLAND, AND E. MANSUR (2012): “What Do Emissions Markets Deliver and to Whom? Evidence from Southern California’s NOx Trading Program,” *American Economic Review*, 102, 965–993.
- HECKMAN, J. J., H. ICHIMURA, AND P. TODD (1998): “Matching as an Econometric Evaluation Estimator,” *The Review of Economic Studies*, 65.
- HECKMAN, J. J., H. ICHIMURA, AND P. E. TODD (1997): “Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme,” *The Review of Economic Studies*, 64, 605–654.
- IPCC (2014): *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, Geneva, Switzerland: Intergovernmental Panel on Climate Change.
- JARAITE, J. AND C. D. MARIA (2016): “Did the EU ETS Make a Difference? An Empirical Assessment Using Lithuanian Firm-Level Data,” *The Energy Journal*, 37, 1–23.
- KLEMETSEN, M. E., K. E. ROSENDAHL, AND A. LUND JAKOBSEN (2016): “The impacts of the EU ETS on Norwegian plants’ environmental and economic performance,” Working Paper Series 03-2016, Norwegian University of Life Sciences, School of Economics and Business.

- LUTZ, B. J. (2016): “Emissions trading and productivity: Firm-level evidence from German manufacturing,” ZEW Discussion Papers 16-067, ZEW - Center for European Economic Research, Mannheim.
- MARIN, G., M. MARINO, AND C. PELLEGRIN (2017): “The Impact of the European Emission Trading Scheme on Multiple Measures of Economic Performance,” *Environmental and Resource Economics*, forthcoming.
- MARTIN, R., M. MUÛLS, AND U. J. WAGNER (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., M. J. MELITZ, AND G. I. P. OTTAVIANO (2014): “Market Size, Competition, and the Product Mix of Exporters,” *American Economic Review*, 104, 495–536.
- MINISTÈRE DE L’ENVIRONNEMENT (2005): “Décret n° 2004-832 du 19 août 2004 pris pour l’application des articles L. 229-5 à L. 229-19 du code de l’environnement et relatif au système d’échange de quotas d’émission de gaz à effet de serre.” .
- PETRICK, S. AND U. J. WAGNER (2014): “The Impact of Carbon Trading on Industry: Evidence from German Manufacturing Firms,” Kiel Working Papers 1912, Kiel Institute for the World Economy.

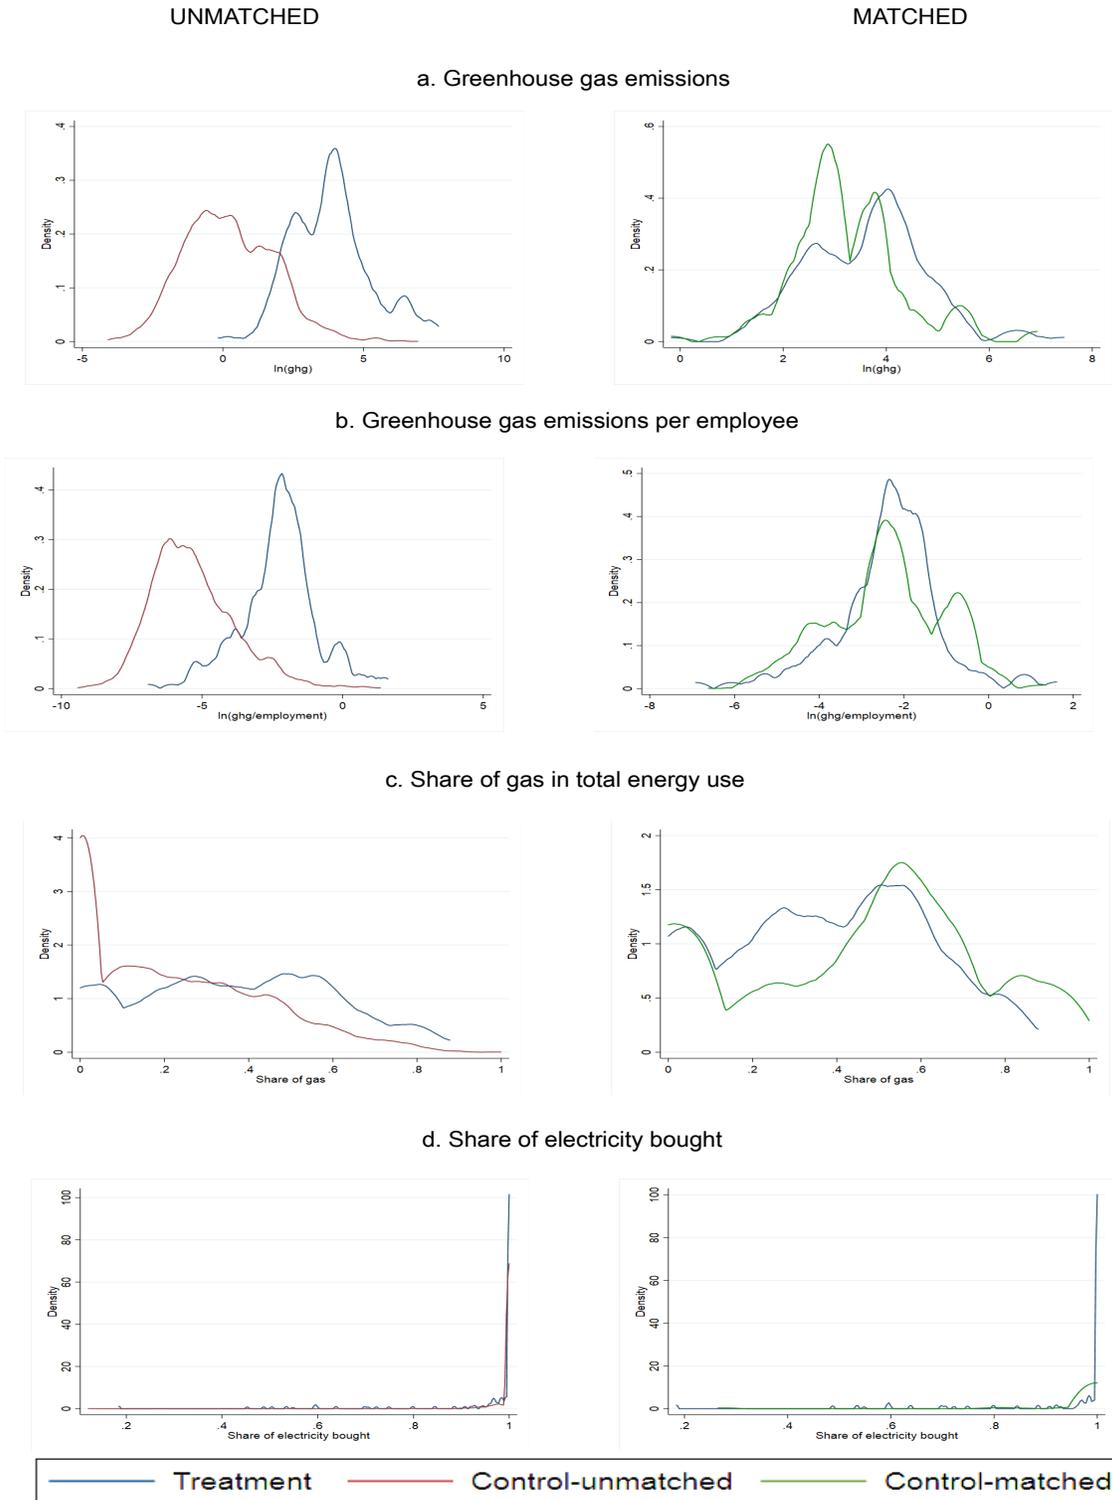
# Online Appendices – Not for Publication

## A Matching Appendix

### A.1 Pre- and Post-Matching Balance

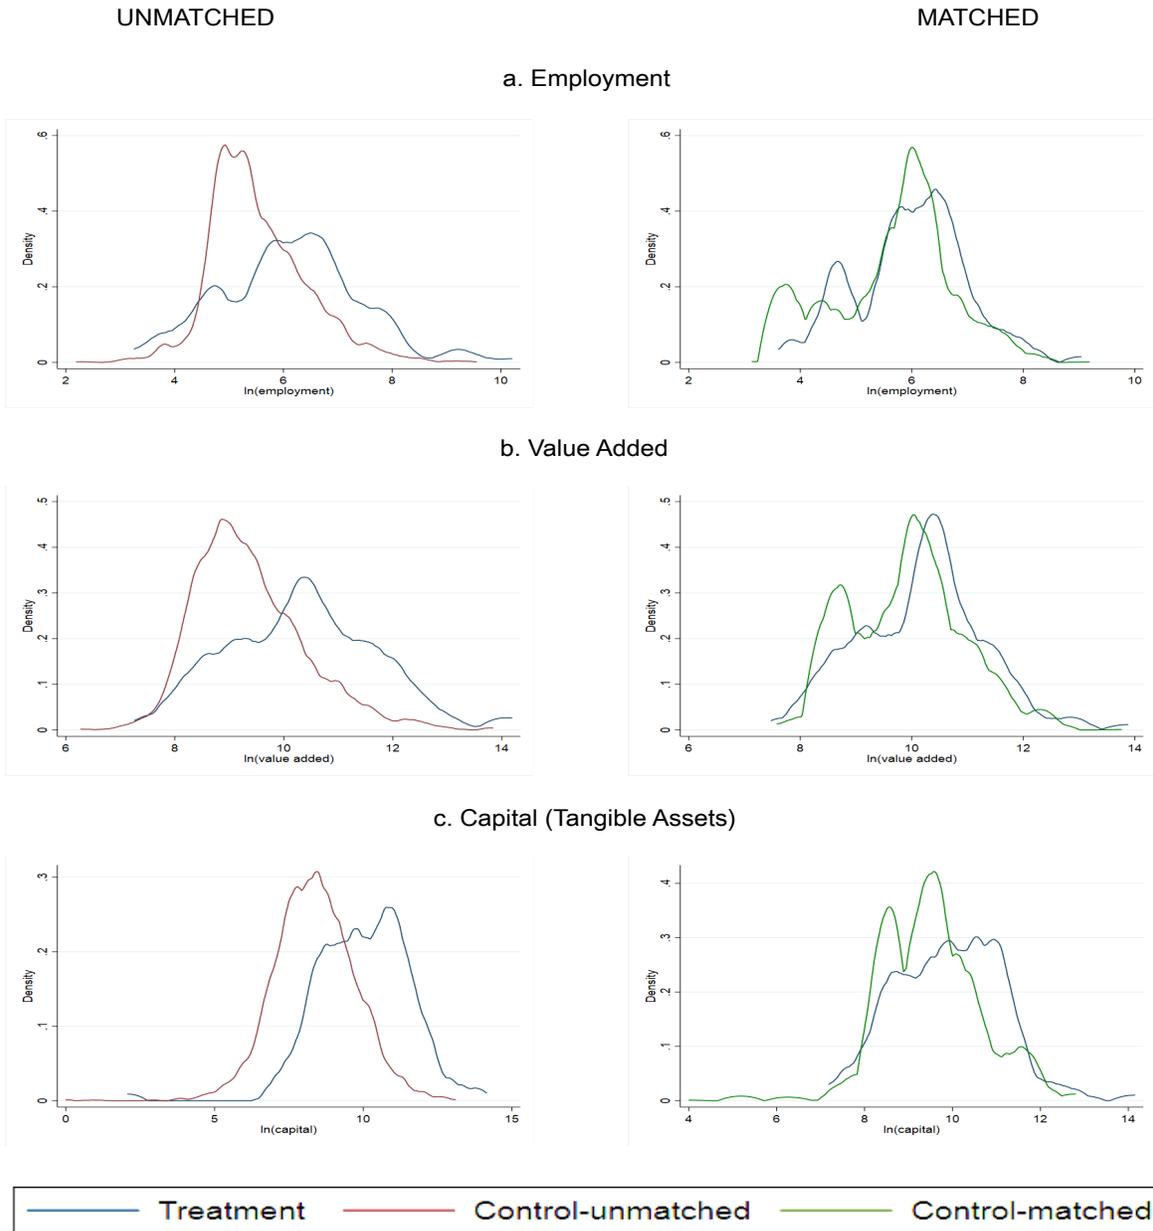
Figures [A.1](#) and [A.2](#) illustrate the density plots of treated vs. untreated firms for different samples. The left column compares both groups of firms before the matching while the right sample compares both groups in the matched sample. They show that the matching approach improves considerably common support.

Figure A.1: Density plots for pre- and post-matching balance - energy



Notes: The figures report the density plots of the variables' values in 2000, our base year for the treated group of firms and the control, untreated, group of firms in the unmatched and matched sample.

Figure A.2: Density plots for pre- and post-matching balance - other



Notes: The figures report the density plots of the variables' values in 2000, our base year for the treated group of firms and the control, untreated, group of firms in the unmatched and matched sample.

## A.2 Alternative Matching Specifications

We here present the regression results for greenhouse gases in Table 2 when applying different matching approaches. Column 1 of Table A.1 is identical to Column 1 of Table 2 with our main result regarding emissions, based on matching on the logarithm of CO<sub>2</sub> in 2000 and exactly on sector. We also impose maximum nearest-neighbor distance and in the number of plants difference. In the other columns of the table, the same matching is applied, except that instead of matching on the logarithm of CO<sub>2</sub>, we also match on other variables. Column 2 shows that matching on both the logarithm of CO<sub>2</sub> and the logarithm of employment, the parallel trends assumption does not hold. Column 3 shows that the results from column 1 are robust to matching also on the percentage of electricity consumed that the firm buys from the grid.

Table A.1: Matching specifications

	(1)	(2)	(3)	(4)	(5)
	Dependent variable is $\Delta \log(\text{CO}_2)$				
PRE	0.0110 (0.0202)	-0.0408** (0.0178)	0.0116 (0.0191)	0.0144 (0.0197)	0.0110 (0.0203)
ANNOUNCEMENT PHASE	0.0655*** (0.0248)	-0.00741 (0.0334)	0.0209 (0.0245)	0.0709*** (0.0253)	0.0667*** (0.0251)
PHASE I	0.000559 (0.0448)	-0.00401 (0.0463)	0.0229 (0.0501)	0.00750 (0.0474)	0.00922 (0.0445)
PHASE II	-0.135** (0.0591)	-0.130** (0.0637)	-0.108* (0.0622)	-0.145** (0.0589)	-0.134** (0.0589)
Observations	1,575	1,660	1,647	1,613	1,567
Matching	ln(CO <sub>2</sub> )	ln(CO <sub>2</sub> ) & ln(Employment)	ln(CO <sub>2</sub> ) & ln(CO <sub>2</sub> /VA)	ln(CO <sub>2</sub> ) & % electricity bought	ln(CO <sub>2</sub> ) & & % Coal share

NOTES: Significance levels are indicated as \* 0.10 \*\* 0.05 \*\*\* 0.01. Distance restrictions between treatment and control firms are imposed at the 95th percentiles. Standard errors are two-way clustered by firm and matching group. Different matching specifications are presented in each column as specified by the line “Matching”

### A.3 Distance Restrictions

We also present in Table A.2 how sensitive the results are to the choice of distance restriction that is placed when matching to the control. Our main specification from Table 2 is reproduced in column two. We find that when allowing for less distance, the significance and size of the reduction in CO<sub>2</sub> emissions in Phase II is reduced.

Table A.2: Distance restrictions

	(1)	(2)	(3)
	Dependent variable is $\Delta \log(\text{CO}_2)$		
PRE	0.0137 (0.0199)	0.0110 (0.0202)	0.00880 (0.0206)
ANNOUNCEMENT PHASE	0.0652*** (0.0246)	0.0655*** (0.0248)	0.0678*** (0.0249)
PHASE I	-0.00125 (0.0455)	0.000559 (0.0448)	0.00750 (0.0443)
PHASE II	-0.138** (0.0577)	-0.135** (0.0591)	-0.106* (0.0581)
Observations	1,620	1,575	1,502
Distance percentile	99th	95th	90th

NOTES: Significance levels are indicated as \* 0.10 \*\* 0.05 \*\*\* 0.01. Distance restrictions between treatment and control firms are imposed at different percentiles. Standard errors are two-way clustered by firm and matching group.

## B Additional Results and Robustness Tests

### B.1 Difference-in-Difference

In this section we present evidence based on the more traditional parametric difference-in-difference approach. Specifically, we estimate the following model, using the full sample of firms,

$$y_{jt} - y_{j2000} = \alpha + \delta \mathbb{1}(ETS_j) + \sum_{\tau} \gamma_{\tau} \mathbb{1}(period_{\tau}) + \sum_{\tau} \beta_{\tau} \mathbb{1}(period_{\tau}) \times \mathbb{1}(ETS_j) + \varepsilon_{jk}$$

where  $y_{jt} - y_{j2000}$  denotes the difference in outcome for firm  $j$  in year  $t$  relative to the year prior to the announcement of the EU ETS, the year 2000. The coefficients of interest are  $\beta_{\tau}$ , where  $\tau = t - 2000$  captures the effect of the EU ETS on regulated firms, relative to unregulated firms, in each year relative to the year 2000.  $\gamma_{\tau}$  captures the effect of the EU ETS on unregulated firms. Standard errors are clustered at the firm-level to account for serial correlation over time within-firms.

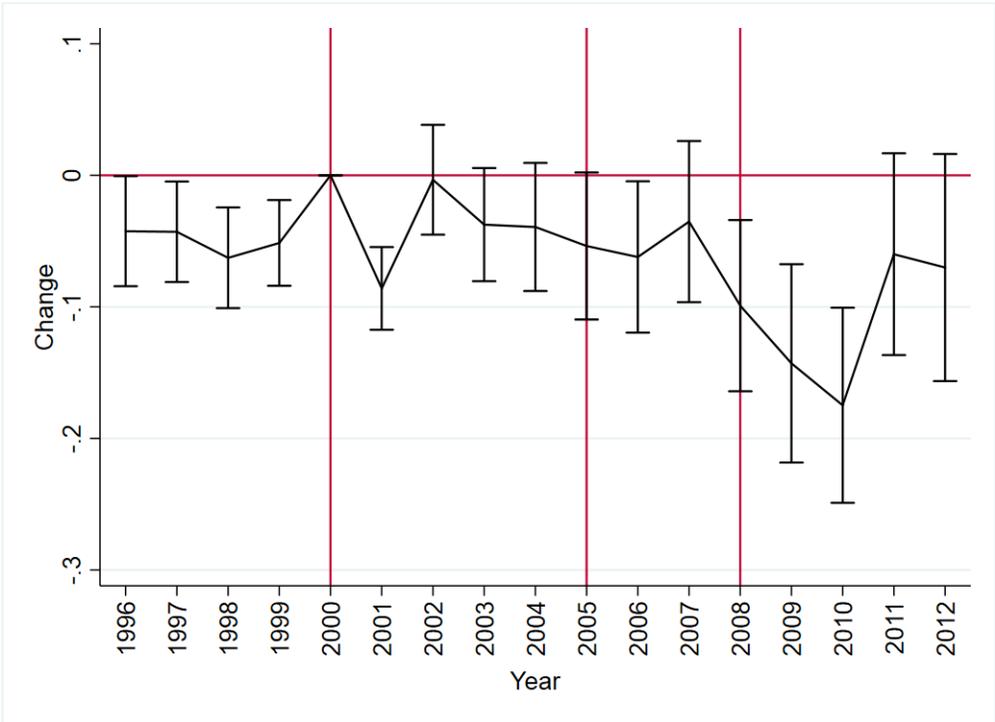
Table B1 reports the average treatment effect on the treated for each period of the EU ETS: the pre-treatment period (1996-2000), the announcement phase (2001-2004), Phase I (2005-2007) and Phase II (2008-2012). We observe that for all four dependent variables, the parallel trends assumption is rejected by the data. Figure B.1 presents graphically the average treatment effect on the treated for each year relative to the year 2000 for CO<sub>2</sub> emissions.

Table B1: The Effect of the EU ETS on Emissions and Economic Outcomes – Parametric Difference-in-Difference

	(1)	(2)	(3)
	$\Delta \log \text{CO}_2$	$\Delta \log \text{Value Added}$	$\Delta \log \text{Employment}$
Pre-Announcement	-0.0263* (0.0135)	0.0310* (0.0181)	0.0527*** (0.0121)
Announcement Phase	-0.0442*** (0.0169)	-0.0141 (0.0208)	-0.0107 (0.0123)
Phase I	-0.0492* (0.0272)	-0.0507 (0.0339)	-0.00515 (0.0211)
Phase II	-0.110*** (0.0322)	-0.0264 (0.0418)	-0.00137 (0.0280)
Observations	38,644	38,644	38,644
Firms	2,871	2,871	2,871

NOTES: Significance levels are indicated as \* 0.10 \*\* 0.05 \*\*\* 0.01. Standard errors are clustered at the firm level.

Figure B.1: Change in CO<sub>2</sub> emissions at treated vs. control firms - Parametric Difference-in-Difference



Notes: The figure represents the differences in  $\log(\text{CO}_2 \text{ emissions})$  between treatment and comparison firms over time, with the same approach as for the results in Table B1 and with 2000 as the base year. Confidence interval bars at 95%.

Table B2: The Effect of the EU ETS on Firm Energy Outcomes

	(1)	(2)	(3)	(4)
	$\Delta$ Coal share	$\Delta$ Gas share	$\Delta$ Elec. share	$\Delta$ Elec. bought sh.
PRE	0.00144 (0.00295)	-0.000210 (0.0123)	0.000599 (0.00333)	-0.0187 (0.0115)
ANNOUNCEMENT PHASE	0.00412 (0.00446)	0.0226 (0.0162)	0.0329*** (0.0125)	0.0151* (0.00857)
PHASE I	-0.0139 (0.00906)	0.0363* (0.0209)	0.000258 (0.0142)	0.0140* (0.00842)
PHASE II	-0.00731 (0.00759)	-0.0171 (0.0313)	0.000184 (0.0191)	0.0254 (0.0157)
Observations	1,575	1,575	1,575	1,573

NOTES: Significance levels are indicated as \* 0.10 \*\* 0.05 \*\*\* 0.01. Distance restrictions between treatment and control plants are imposed at the 95th percentile. Standard errors are two-way clustered by firm and matching group.