The Causal Effect of the European Union Emissions Trading Scheme: Evidence from French Manufacturing Plants*

Ulrich J. Wagner
Universidad Carlos III

Mirabelle Muûls
Imperial College London

Ralf Martin
Imperial College London

Jonathan Colmer
London School of Economics

Abstract

How do firms respond to regulation and market incentives? Using comprehensive plant-level data for around 9,500 French manufacturing firms, this paper explores the economic and environmental response of plants to the European Union Emissions Trading Scheme (EU ETS) – the EU's flagship climate policy. Our results suggest that ETS-regulated manufacturing plants in France reduced emissions by an average of 15.7%. The most marked reduction in emissions is observed during Phase II (2008-2013). In terms of economic outcomes, we find a statistically significant reduction in employment (10.4% in Phase II). We aim to understand whether these reductions are real global reductions in emissions or whether they are merely the result of carbon leakage. Leakage may occur because output shifts to unregulated plants, either within the EU or abroad. It may also occur because regulated plants outsource carbon intensive parts of the production process. We make some progress in providing evidence on this by looking at the fuel mix of plants. About half of the reduction in emissions can be accounted for by an increase in the share of gas, which is less carbon intensive than coal and oil. This is both consistent with carbon leakage, because of outsourcing, and real emissions reductions due to a technological change. Ongoing research will shed further light on this in the future. We also examine if there is leakage within firms by looking at firms with both, unregulated and regulated plants. However, we find no evidence in support of this.

*Still preliminary. Wagner: Universidad Carlos III de Madrid, Department of Economics, Calle Madrid, 126 28903 Getafe, Spain. Phone: +34 (0) 916 24 8488, email: uwagner@eco.uc3m.es. Muûls: Grantham Institute for Climate Change and Imperial College Business School, Imperial College London, South Kensington Campus, London SW7 2AZ and Centre for Economic Performance, London School of Economics, Houghton Street, London WC2A 2AE. E-mail: m.muuls@imperial.ac.uk. Martin: Imperial College Business School, Imperial College London, South Kensington Campus, London SW7 2AZ and Centre for Economic Performance, London School of Economics, Houghton Street, London WC2A 2AE. E-mail: r.martin@lse.ac.uk. Colmer: Centre for Economic Performance and Grantham Research Institute, London School of Economics and Political Science, Houghton Street, London WC2A 2AE. E-mail: j.m.colmer@lse.ac.uk. This project was gratefully supported by research funding from the Economic and Social Research Council (grant ES/J006742/1), the ESRC Centre for Climate Change Economics and Policy and the Grantham Foundation is gratefully acknowledged.

1 Introduction

In the presence of market failure, the role of government intervention is naturally debated. If such intervention is to happen, one must understand how economic agents respond to regulation and market incentives. By understanding this behaviour, interventions can be designed to reduce their social cost.

One of the clearest examples of market failure in the 21st Century is climate change. To the degree that climate change is driven by increases in the atmospheric concentration of greenhouse gases – a position that is strongly supported by the physical science literature (IPCC, 2014) –, it is reasonable to argue that *"climate change...is the greatest and widest-ranging market failure ever seen"* (Stern, 2006).

The European Union has been a central player in the global efforts to curb greenhouse gas (GHG) emissions and mitigate climate change. In 2005, the EU launched Phase I of the Emissions Trading System (EU ETS), the first international system for trading GHG emissions. Following a three-year pilot period, Phase II of the EU ETS was launched in 2008. Across its 28 Member States (MS) and in three European Economic Area (EEA) states, the EU ETS covers large plants from CO_2–emission intensive industrial sectors and all other combustion activities with a rated thermal input above 20MWh, covering around 45% of the EU's emissions (5% of global emissions). During the first two years, this scheme included approximately 10,600 industrial plants from 25 MS. There are now 31 countries and more than 11,400 plants taking part in the Scheme.

This paper sets out to understand the impact of the EU ETS on the economic and environmental performance of regulated installations, through the use of comprehensive plant-level data between 1992 and 2010. Compared to studies at the aggregate level, which have to identify effects by projecting a baseline into the future, we are able to account for idiosyncratic shocks affecting the economy as a whole and specific sectors over time. Furthermore, the use of plant level data allows us to compare the performance of plants that are regulated under the EU ETS to plants with similar characteristics that are not regulated.

The central outcome of interest for a policy such as the EU ETS are CO_2 emissions. The only source for representative emissions data for both EU ETS and non–EU ETS plants are confidential business surveys maintained by government statistical agencies. Access to these datasets is restricted and subject to disclosure control. This explains why studies of this kind haven't been undertaken to date. This paper uses administrative panel data on French manufacturing plants to shed light on this issue.

There is also considerable concern regarding the impact of the EU ETS (and environmental regulation more broadly) on economic outcomes, specifically impacts on the competitiveness, productivity, and employment of regulated installations. On the one hand, because of the unilateral nature of the EU ETS and the integration of economic activities in global markets, lobby groups and policy makers have long argued that environmental regulation may undermine the competitiveness of regulated firms and "kill jobs". On the other hand, more stringent regulation may enhance productivity as less productive firms exit the market (Porter, 1991). One reason that debates over the economic consequences of environmental regulation have raged on for so long is the paucity of conclusive empirical evidence (Greenstone, List, and Syverson, 2012). Official datasets provide one of the most reliable sources for this kind of information and we explore these outcomes as well.

Our results suggest that ETS-regulated manufacturing plants reduced emissions by an average of 15.7%, compared to non-ETS plants, since the implementation of the policy in 2005. The most marked reduction occurs following the onset of Phase II of the EU ETS in 2008, though there a some signs of reductions during Phase I (2005-2007). While this is promising in terms of a climate policy objective, we also find a 10.4% reduction in employment for regulated plants during phase II – a concern, potentially, for stakeholders worried about the economic consequences of carbon markets.

The consequences of our results are as follows: either the EU ETS has been very successful in reducing emissions, but with significant reductions in employment, or it has resulted in carbon leakage where emissions, as well as employment and output, are shifting from regulated to unregulated facilities, either abroad or within France.

There are a number of possible channels and mechanisms through which carbon leakage could occur. First, regulated firms might move production to unregulated plants (at home or abroad). Secondly, regulated firms might become less competitive and lose market shares to firms in jurisdictions with less stringent carbon regulation. Thirdly, regulated firms and facilities might outsource their more carbon intensive activities to unregulated plants. This case would be consistent with leakage; however, there would be little or no change in the final output of regulated firms.

In the absence of global firm-level data, as well as appropriate control groups, we can not fully distinguish between these different channels. Neither do we have the power to reject the hypothesis of no leakage. Nevertheless, we can make progress on this issue in a number of ways.

First, we examine whether there is any evidence of changes in the fuel mix used by regulated plants. If leakage is occurring due to channels one and two, we should expect a reduction in regulated facility output with little impact on the facility level fuel mix. In turn, evidence of changes in fuel mix are supportive of leakage channel three. However, they could also indicate that firms achieve real emission reductions by changes in the production technology.

Initial results suggest that the EU ETS leads to significant and economically meaningful changes in the fuel mix: we find reductions in the share of coal and oil used in regulated plants by 4.77% and 4.92% respectively during Phase II. This reduction in the share of coal and oil was offset by an increase in the use of natural gas, whose share increased by 7.66%. This corresponds to a reduction in the carbon intensity of the fuel mix as the carbon content of natural gas, while positive, is substantially less than coal and oil. A simple back-of-the-envelope calculation suggests that fuel mix changes account for around 50% (7 percentage points) of the observed emission reductions.

To refine our understanding of whether this reduction reflects outsourcing or is the result of technological change it would be desirable to have information on intermediate input intensity. If outsourcing is a major driver, we should see an increase in intermedates' shares. If firms outsource from unregulated countries outside the EU we should also see an increase in firm-level import intensity. In the current version of the paper we cannot provide this. However, we are in process of gaining access to this data from the relevant French agencies and we will examine this in a future version of the paper.

Secondly, having facility level data, we can explore if there is any evidence of within firm leakage for firms with both unregulated and regulated facilities. Arguably - all else equal - we would expect that it would be easier for such firms to shift emissions to unregulated plants as they are incurring less transaction costs than firms who have no preexisting links with unregulated facilities. However, we do not find any evidence for such effects.

Hence, based on our current results we have no reason to believe that the emission reductions we find are the consequence of carbon leakage. However, we equally cannot reject that no leakage occurs. With data on intermediates and revenues we might be able provide some lower bounds on global emission reductions. For instance, suppose we cannot find evidence to support theoutsourcing of carbon intensive production processes. In this case the 7 percentage point reduction in emissions, due to fuel mix changes,is likely driven by technology changes, thereby providing a lower bound for emission reductions.

A further point of concern is the timing of the emission reductions which appears primarily in Phase II of the ETS after 2007. Emissions trading provides incentives to firms to reduce emissions primarily due the actual or opportunity cost of emission permits. However, during the second phase of the ETS, the price of carbon fell due to a combination of forces, such as the financial crisis and a strong contraction of permit demand. This begs the question as to why regulated plants would reduce emissions any more than unregulated plants. In section 6 we introduce a simple theory of the firm that provides an explanation for this puzzle. Most stakeholders expect that prices will increase in the future (Martin et al., 2013a). Given that investments in abatement capital today are durable and will consequently last until prices increase, we might expect emissions reductions today even if current prices are low. This model implies that an increase in investment is a sufficient statistic for emission reductions. This result also helps to distinguish whether the negative employment effects are due to leakage or complementarities with carbon-intensive capital. In the latter case, the ETS should result in an increase in investment. In the former case, we may expect a reduction in investment, and perhaps an increase in intermediary goods. The theory also has implications for heterogenous treatment effects. It implies that expected future price increases can have a positive effect on current emissions if firms decide to exit in the future. Fixed investments will be used for a shorter amount of time if firms intend to exit, and so they will reduce investments, increasing emissions. This result arises from the durability of emission-saving investments. Consequently, we might expect to see an increase in emissions for firms where demand is inelastic and the substitutability of abatement technology with "dirty" capital is low. Finally, firms with a younger capital stock are less likely to respond to changed future price expectations and so may not reduce emissions during phase II. Understanding the empirical relevance of these effects is an additional focus of our current research efforts.

The remainder of the paper is structured as follows: the next section reviews the existing economics literature on the EU ETS; section 3 describes the two datasets used in this paper; in Section 4 we discuss the econometric approach adopted. Section 5 presents our findings on the effect of the EU ETS on environmental and economic performance of regulated plants; section 6 introduces a model motivated by our findings and provides a discussion of its implications with a focus on our current research efforts; section 7 concludes.

2 Existing Literature

There are a number of studies on the EU ETS. While the financial economics literature has analysed the workings of the allowance market (e.g. Alberola et al. (2009), Trotignon and Delbosc (2008)), one policy variable that is heavily debated at both the academic and policy level is the practice of allocating emission permits. Ellerman et al. (2007) give a detailed account of the development of the NAPs under Phase I in 10 European countries. Important dynamic effects of a permit allocation scheme that derive from the treatment of plants that close and of new entrants are analysed by Åhman et al. (2007). Previous research has also shown that these variations in the allocations of different firms regulated by the EU ETS are important for innovation (Martin, Muûls, and Wagner, 2013a).

A natural consequence of the implementation of the EU ETS is that firms in the ETS face higher carbon prices than their competitors outside the EU who are not subject to comparable regulation. This has led to worries about a possible loss of competitiveness of European industry. A widespread approach to assessing these effects has been to calibrate computable general equilibrium (CGE) models that are capable of predicting the consequences of differential carbon pricing across regions and the resulting carbon leakage (see the survey by Paltsev, 2001). Another strand of research conducts ex ante analyses of the competitiveness effects of the EU ETS, based on simulations and with particular attention to sectoral detail. Existing studies (such as McKinsey and Ecofys (2006) or Reinaud, 2005) conclude that competitiveness effects are moderate as long as permit allocation is free of charge. As a larger share of permits will be auctioned in the future, the most energy-intensive industries will be at risk of a competitiveness loss. Grubb, Brewer, Sato, Heilmayr, and Fazekas (2009) argue that such detrimental competitiveness impacts are limited to a small number of industry sectors.

Very few ex post evaluation studies of the competitiveness effects of the EU ETS have been completed to date. Martin et al. (2013b) review this literature in detail. Demailly and Quirion (2008) and Anger and Oberndorfer (2008) study the impact of the EU ETS on production and profitability respectively for the sector of iron and steel industry in a sample of German firms. The former study finds modest competitiveness losses and the latter no significant impact on revenues or employment of the firms under regulation. Recent research has also established that the existing allocation rules by the EU are not optimal and could be modified to achieve greater emissions reductions as well as a lower risk of carbon leakage and loss of competitiveness (Martin, Muûls, de Preux, and Wagner, 2014b,c). Aside from that, researchers have been keen to evaluate the effectiveness of the EU ETS at incentivizing carbon abatement. Ellerman and Buchner (2008) and work by Delarue and D'haeseleer (2007) combine the Community Independent Transaction Log (CITL) and carbon price data to perform counterfactual simulations at the sector level. Ellerman et al. (2010) also use a macro approach as well as data on the electric utility sector to study abatement. All three show the overall success of the market-based instruments on emissions abatement, even though the emission caps in Phase I turned out not to be binding. Emission caps in Phase II have been more ambitious so that the risk of over-allocation has been substantially reduced (Ellerman and Joskow, 2008). [1]

Petrick and Wagner (2014) use administrative data on German manufacturing firms to estimate the impact of the EU ETS using nearest-neighbor matching. They find a significant reduction in carbon emissions during the first half of phase II (the last year included is 2010) in the order of 25% at treated firms. While they find no effect on employment, they find moderate but statistically significant increases in the value of output and exports during the second phase – interpreted as the pass-through of carbon prices onto consumers. In contrast to their firm-level analysis, this paper analyses the policy impact at the plant level - the finest level of disaggregation available to researchers and the level at which treatment is allocated under the ETS when a plant has only one installation or if all plant installations are regulated.

In summary, much of the econometric evidence on the effects of the EU ETS so far is limited in scope and based primarily on sector-level analysis. While such an approach gives a first insight on the effects of the policy, it does not provide conclusive evidence to disentangle the causal effect of the EU ETS from other concurrent events. Furthermore, while previous efforts have been made to identify the magnitude of the ETS on environmental and economic performance, there have been minimal efforts to understand the driving forces behind these effects. We will address these problems by making use of longitudinal observations at the level of the individual business enterprise, for both ETS and non-ETS plants, to study the effects of the policy on energy use and economic performance and exploit the granularity and richness of the data to better understand the driving forces underlying the estimated effects.

3 Data

3.1 Annual Survey of Industrial Energy Consumption

The EACEI (L'Enquete Annuelle sur les Consommations d'Energie dans l'Industrie) is a survey

[1] As far as other climate policy measures are concerned, it has been shown that the Climate Change Levy for UK manufacturing firms had no significant impacts on employment, gross output or total factor productivity and that a substantial tax discount granted to some firms prevented further cuts in energy use (Martin, de Preux, and Wagner, 2014a).

conducted annually in France.[2][3] It provides quantities and values of energy consumed by energy type as well as the different usages of each type of energy. Other variables provided in the survey include employment, geographical location and sectoral classification. Information for the following energy types is requested from the surveyed firms: electricity (bought, auto-produced and resold), vapour, 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. Electricity usages include: driving force, thermal uses, other uses (including electrolysis). For other types of energy, the survey distinguishes between: manufacturing, electricity production, raw materials, heating and other purposes.

Until 2007, firms covered by the survey included those in sectors 12 to 37[4] according to the NAF rev.1 classification, equivalent at the two-digit level to the NACE rev.1. In the most recent years, about 12,000 establishments are part of the sample: all industrial establishments employing 20 employees or more in the most energy consuming sectors (23.32Z, 23.51Z and 23.52Z); all establishments with more than ten employees in sector 20.11Z (manufacturing of industrial gases); all establishments with more than 250 employees on the 31st of december of that year; a sample of establishments with employment between 20 and 249 employees in sectors that are not energy intensive. The sampling frame includes all French manufacturing enterprises and establishments. 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. The response rate is close to 90%.

3.2 Annual Business Survey

The EAE (L'Enquete Annuelle des Entreprise) collects balance-sheet data at the firm level on turnover, employment, capital, and aggregate wages, as well as information about firm location and industry classification. The data is less detailed at the plant level. The data are available for the period 1992 to 2010 for firms with more than 20 employees and all the plants of those firms.

3.3 The European Union Transaction Log

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 interlink with the Community Independent Transaction Log (or CITL), operated by the Commission, which records and checks every transaction. It is now centralised as the

[2] Until 2008 it was conducted by the Service des études et des statistiques industrielles (SESSI), the ministry in charge of manufacturing. Since 2009 the EACEI has been conducted by the Statistical direction of INSEE, the National Institute of Statistics and Economic Studies.

[3] All plant-level data used here was provided for research purposes by authorization of the Comité du Secret Statistique.

[4] except sectors 15, 16, 20.1A, 22.1 and 23

European Union Transaction Log (or EUTL). Each of the approximately 12,000 installations has an ”operator holding account” in its national registry, into which its own allowances are issued. Any individual or organisation wishing to participate in the market is able to open up their own “person holding account” in any of the registries. The CITL's web pages makes publicly available contact details for each account, the number of allowances allocated under the “national allocation plan” and the compliance position of each plant. Records of other types of transactions are made available after a period of five years has passed.

In France, the national registry is managed by the Caisse des Dépôts. Their website provides additional information, including a link between the permit identifier (GIDIC) from the national registry and the SIREN (Système d'Identification du Répertoire des Entreprises) identifier from the INSEE. 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, which allows a quasi-perfect matching of the two databases through the SIREN and postcode identifiers. When multiple plants of a EU ETS participating enterprise share the same postcode, name matching and latitude-longitude information available in the CITL was used to ensure a better match.

3.4 Descriptive statistics

The resulting dataset includes 11,575 plants within 9,494 firms. 368 plants within 279 firms are part of the EU ETS. Table 1 presents descriptive statistics of the different variables used in our econometric analysis. Table 2 demonstrates how ETS and non-ETS plants differ within these dimensions. Table 3 presents descriptive statistics of the different variables used in our econometric analysis for ETS plants, based on whether they belong to firms with multiple plants that only have ETS plants; to firms with both ETS and non-ETS plants; or whether they belong to a single-plant ETS firm. We observe that 29% of ETS plants are part of firms with non-ETS facilities. Within this subset of firms, 62% of the plants are not regulated by the ETS. Overall, we see that ETS plants are on average larger, emit more GHG emissions and are more carbon intensive. The following section describes how we construct a representative control group to compare regulated and non-regulated plants given the observed differences in covariates.

4 Research Design

In this study, we exploit variation in the selection criterion by which plants are required to join the EU ETS. Building on the potential outcomes framework commonly used in the program evaluation literature, we propose that plants can be in one of two states: either part of the market-based EU ETS, or prevailing in a state of business as usual.[5]

Let $ETS_i = 1$, if plant i is a member of the EU ETS and is therefore part of the “treatment group”. Let $ETS_i = 0$ if plant i is not part of the EU ETS and is therefore part of the

[5]See Holland (1986) for a deeper discussion of causal inference, the potential outcomes framework, and its history.

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 plant 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):

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

where α_{ATT} measures the average effect of the EU ETS on the outcome variable of interest, namely plant level emissions. Emissions at both treated and untreated plants are observed for several years prior to the announcement of the EU ETS in 2000, prior to its implementation in 2005 and for Phases I and II of the scheme.[6]

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

The crudest and most naive estimate of the α_{ATT} is obtained by computing an unconditional difference-in-difference. However, one of the major constraints in estimating the causal effect of the EU ETS is constructing a suitable counterfactual with which to compare EU ETS plants. If there are significant differences between the characteristics of ETS and non-ETS plants that are correlated to plant-level emissions dynamics, then estimates of the causal effect of the EU ETS on emissions and economic outcomes may be seriously biased (see Heckman, Ichimura and Todd, 1998).

$$\alpha_{ATT}^{biased} = \mathbb{E}[\Delta Y_i(1)|ETS_i = 1] + \mathbb{E}[Y_{0i}|ETS_i = 1] - \mathbb{E}[Y_{0i}|ETS_i = 0].$$

To reduce the bias introduced by observable differences between ETS and non-ETS plants, we employ a semi-parametric conditioning strategy that matches treated plants to non-treated plants based on observable characteristics in the spirit of Heckman, Ichimura, and Todd (1997; 1998), Blundell et al. (2001) and Abadie (2005).

4.1 A Semi-parametric conditioning strategy

As discussed, a simple comparison of ETS plants with non-ETS plants, even after controlling for observables, may still result in bias, attributing some of the changes between the outcome variables to the EU ETS when they are really the result of other systematic differences between treatment and control plants. Matching estimators, an extension of standard regression approaches, can help us to reduce this bias, while avoiding the parametric assumptions about the relationship between the outcome variable and the control variables in X_i. Our approach follows Heckman et al. (1997, 1998), who implement the following generalized difference-in-difference estimator:

[6] Emissions are converted from fuel use from the EACEI database using standardized conversion factors to be used by EU ETS participants in France provided by the ADEME agency.

$$\alpha_{ATT} = \mathbb{E}[Y_{i1}(1) - Y_{i0}(1)|X_i, ETS_i = 1]$$
$$= \frac{1}{N_1}\sum_{j\in I_1}\{(Y_{jt_1}(1) - Y_{jt_0}(0)) - \sum_{k\in I_0}\omega_{jk}(Y_{kt_1}(0) - Y_{kt_0}(0))\}$$

where I_1 denotes the set of EU ETS plants, I_0 the set of non-ETS plants, and N_1 the number of participating plants in the treatment group. The treated plants are indexed by j; the control plants are indexed by k. The weight placed on a non-ETS plant when constructing the counterfactual estimate for EU ETS plant j is ω_{jk}. These weights can be calculated using any matching approach. This imposes the same distribution of covariates for ETS and non-ETS plants (Rosenbaum and Rubin, 1983).

A final assumption, necessary to rule out spillovers and general equilibrium effects, is that potential outcomes at one plant are independent of the treatment status of other plants. This is known as the stable unit treatment value assumption (SUTVA; Rubin, 1980). While it is straightforward to demonstrate that the common support condition and parallel pre-trends assumption are satisfied, other assumptions, such as SUTVA are not testable; however, we can implement indirect tests to evaluate whether these assumptions are plausible. This is important as it might be the case that within a firm, emissions are reallocated from ETS plants to non-ETS plants. A further concern are possible general equilibrium effects between treatment and control plants. If the ETS affects the productivity of larger firms, non-regulated plants may increase their market share, increasing output and emissions. This may inflate our estimate of the average treatment effect. By aggregating our analysis to the firm level (forthcoming) we can check for potential spillovers within-firm, although we are unable to rule out spillovers between firms. Understanding the effects of regulation and market incentives on market power and competition is an interesting area for future research. For all plant-level regressions the standard errors are clustered at the firm level to account for any within-firm correlation across plants.

5 Results

In this section we report the results from the semi-parametric difference-in-difference approach. For each EU ETS plant and for a fixed base-year, we identify the nearest neighbour match based on a propensity score.[7] Given the difficulties associated with size differences between plants, we estimate the propensity score using the carbon intensity of each plant in the year 1999, the announcement year of the EU ETS. We also match each plant exactly by sector at the two-digit level.[8] Balance tests are reported in the appendix.

[7]Results, available on request, are robust to nearest neighbour matching based on Euclidean distance or mahalanobis weighting.

[8]Results, available upon request, are broadly robust to matching on additional covariates. To some degree we are limited by the number of covariates available to match on, namely GHG emissions, employment, and carbon factor intensity. Results are also robust to matching on carbon factor intensity and matching exactly

We begin by examining the effects of the EU ETS at the plant level. We then test the robustness of these results at the firm level and decompose any effects based on the structure of the firm in order to identify the presence, if any, of strategic behaviour and leakage.

All observations prior to the year 1999 are defined as the pre-announcement phase, a test of the parallel trends assumption. Between 1999 and 2004 we define the announcement phase, to capture any strategic or expectations-related behavior by firms associated with the EU ETS before its implementation. Phase I is defined as the period covering 2005-2007 and Phase II is defined as the period covering 2008-2010.[9] The coefficient in table 4 reports the average of each of these Phases – the average treatment effect on the treated. The average treatment effect on the treated for each year is also reported graphically after each table.

5.1 Plant Level Results

Table 4 presents the baseline results, applying the semi-parametric difference-in-difference strategy to plant-level manufacturing data in France. Each panel reports the average treatment effect of the EU ETS on treated facilities compared to the announcement of the EU ETS in 2000 for a different outcome variable: GHG emissions, Employment, and Carbon Intensity (all in logs).[10]

In panel A, we report the effect of the EU ETS on the GHG emissions of ETS plants compared to non-ETS plants. Crucially, we note that during the pre-announcement period there appears to be no significant difference between ETS and non-ETS facilities. In both the announcement phase and Phase I, we observe no significant differences in emissions between ETS and non-ETS facilities compared to the baseline year 1999. In Phase II we find that the EU ETS has had a significant effect on plant-level emissions, observing a reduction in GHG emissions ranging from 13.3-15.7%, with the effect size increasing as we restrict the maximum distance between matching partners. This is the net effect of real emission reductions and any carbon leakage.

Panel B reports the effects of the EU ETS on employment within ETS plants compared to non-ETS plants. Interestingly, we observe a reduction in employment during the announcement phase and also in Phase II, indicating that the EU ETS may have had an impact on the competitiveness of ETS plants. We observe that, during the announcement of the EU ETS, employment fell by 4.4%, while during Phase II employment fell by 10.4%. The consequences of these employment effects must be carefully considered. Even if the reduction in employment is not associated with carbon leakage, it is unclear whether the EU ETS increased unemployment, or whether employees can easily gain work locally in non-ETS plants. To the degree that the

within sector at the three-digit level, as well as matching exactly within sector and geographic region. The extension of this approach dramatically cuts our sample size based on the likelihood of finding a match. Our estimates fall below conventional significance levels when when we match on Carbon Intensity, Employment, and GHG emissions, while matching exactly within sector measured at the three-digit level with a sample that is 75% of our baseline sample.

[9]Phase II carried on until 2012, but the confidential data for French firms is only currently available until 2010. We are waiting on the delivery of new data up until the end of Phase II.

[10]Results, available on request, are robust when the outcome variables are measured in levels. However, the sample size when imposing the strictest reported caliper (10%) is 67% of the sample when measured in logs.

ETS had effects on the market share of ETS-regulated facilities, there may be an increase in employment opportunities in unregulated plants. This would introduce an upward bias to our estimates of the effect of the ETS on regulated plants, as an increase in employment or emissions by control plants would exacerbate the difference between regulated and unregulated plants. This would be in effect, a violation of the SUTVA assumption; however, the existence and relevance of such spillover effects are interesting in and of themselves.

The EU ETS appears to have had no effect during Phase I on either employment or emissions, indicating that plants may have had little incentive to reduce emissions. This may be the result of plants not being able to bank credits between Phases at the time, as well as an overly generous permit allocation during phase I, which resulted in a plummeting of the market price in April 2006 (see also Bushnell et al., 2013), thus reducing the incentive to make investments to reduce emissions.

Unfortunately, due to data limitations, we are unable to examine the effect of the EU ETS on the carbon intensity of production at the plant level, defined as emissions divided by output. However, we are able to get some measure of carbon intensity by examining the effect of the EU ETS on emissions divided by employment. Panel C reports that the EU ETS had little effect on the carbon factor intensity of production with no robust effects during the main phases of the ETS. The absence of significant effects during the main phases of the EU ETS indicates that the reductions in emissions are being offset by the reductions in employment. Of interest, we observe that the carbon intensity of production increased during the announcement phase, driven by a reduction in employment at this time. Future work aims to better understand the interpretation of this effect through an examination of strategic behaviour by plants in advance of permit allocations. The incentives for firms to respond in advance of the ETS is likely to depend on their awareness of the capacity thresholds, in addition to their position compared to the threshold. Firms that are close to the threshold have an incentive to reduce emissions and capacity in order to avoid, or reduce the likelihood of, regulation. By contrast, firms that are well above the capacity threshold have an incentive to increase capacity and emissions in order to gain as many permits as possible when the allocation is made – a ratchet effect. Understanding the heterogeneity associated with such behaviour is of great interest, though likely beyond the scope of this research design.

Fundamentally, it is important to understand the degree to which these results are the result of global emissions reductions, as opposed to carbon leakage. As discussed in the introduction, carbon leakage can occur within or between firms and within the EU/France or between the EU and unregulated countries further afield.

At face value, the results in this section have the following interpretations: either the ETS has been very successful in reducing emissions, but at great cost to firms and workers (depending on the likelihood of finding reemployment), or it has resulted in a reallocation of economic activity, at least in part, thus mitigating its effectiveness in reducing emissions – a *trade-off* scheme. The remainder of this section aims to better understand how effective the ETS has been in reducing emissions, and the consequences for employment reductions.

5.2 Disentangling the Channels

5.2.1 The Composition of Emissions Across Fuel Types

As discussed above, changes in the fuel mix of emissions can point either to outsourcing – a carbon leakage channel – or to technological change, that potentially leads to global reductions in emissions.

Table 5 reports the impact of the EU ETS on the role of coal, oil, gas and steam as a share of each plant's fuel mix. We observe substantial adjustments in the composition of the fuel mix. Regulated plants appear to reduce their share of coal and oil by close to 5% while increasing their share of natural gas (7.6%). This results in a reduction in the carbon intensity of production. While natural gas increases carbon emissions, the carbon content of natural gas is substantially less than that of coal and oil. Consequently, a shift in fuel use from coal and oil to natural gas can result in a substantial reduction in plant-level emissions without the need to reduce output. A simple back-of-the-envelope calculation combining the coefficients from table 5 with the carbon content of emissions provides an estimate of emission reductions associated due to fuel mix changes. We estimate that this accounts for nearly 50% of the observed emissions reductions, an average reduction in emissions of 7 percentage points in regulated plants compared to non-regulated plants.[11] In further research we will explore if this reduction is due to out-sourcing or due to technology changes. In the latter case this could interpreted as a lower bound estimate of global emission reductions.

5.2.2 The Composition of the Firm

In this section we explore if there is any evidence of within firm leakage. We examine two sets of results in particular: Firstly, heterogeneity in the ETS response of multiplant firms with both ETS and non ETS plants versus firms where all plants are regulated by the ETS. Secondly, we look at the response of non-ETS plants in ETS firms.

Heterogeneity There are several reasons to expect that plants in firms with both regulated and unregulated facilities respond differently to ETS participation if leakage is relevant. Unfortunately, not all reasons suggest that the difference goes into the same direction. On the one hand, we would expect that firms with an additional channel to respond to the ETS regulation - i.e., leakage between plants of the same firm - respond more strongly for a given level of output (a subsititution effect). On the other hand, the opportunity to shift emissions between plants provides such firms a cost advantage so that they are likely to contract output by less than

[11] The carbon content of fuel types is estimated from a simple within-plant fixed effects regression of logGHG on each of the fuel types and the electricity share, omitting steam as the baseline category because its emissions content is zero. The calculation combines these coefficients with the coefficients from table 5 as follows: Coal effect on GHG emissions = -0.0477 $\times$2.169 = -0.103 ; Oil effect on GHG emissions = -0.0492 $\times$1.909 = -0.093 ; Gas effect on GHG emissions = 0.0766 $\times$ = 1.689 = 0.129 . The sum of these three effects provides an estimate of the share of emission reductions associated with a reduction in the carbon intensity of production: -0.103 + -0.093 + 0.129 = -0.067.

firms where all facilities are covered by the ETS. Consequently, we may expect emissiosn to fall by less (an output effect). A priori the net effect is ambigious. Of course there might also be heterogeneity simply because of other differences between firms that have both regulated and unregulated facilities and firms with only regulated facilities. For instance they might be active in different sectors with fundamentally different production technologies.

Non-ETS plants in ETS firms A less ambigious sign of within firm leakage concerns the effect of the EU ETS on non-ETS plants in ETS firms. Compared to suitabale control plants in non-ETS firms we should find an increase in emissions as well as output and employment if within firm leakage is an issue.

Table 6 reports our results. In column (1) we report our baseline estimate from table 4 to provide a comparison for results. In column (2) we report the impact of the EU ETS on ETS plants that are part of firms that contain non-ETS plants compared to all non-ETS plants. Contrary to the idea of a stronger substitution effect for part-ETS firms, but in line with a weaker output effect, we observe a smaller coefficient size compared to the effect across all plants equal to around an 11% reduction in Phase II. Note that we include as potential controls non-ETS plants in ETS firms. If we are primarily concerned with the substitution effect, this should yet again increase our coefficient estimates.[12] We explore this in column (3) by relying only on plants in non-ETS firms as control group. The coefficient estimate is no longer significant at the 10% level. However, this is by and large due to an increase in noise rather than a smaller coefficient estimate.

Column (4) reports the results for firms that have no opportunity to reallocate emissions within-firm, i.e., single-plant ETS firms and multi-plant ETS firms with only ETS plants. We observe that, during Phase II, ETS plants reduce emissions by 21%, compared to non-ETS plants. The magnitude of this coefficient is substantially larger than the magnitude of the coefficients in columns (2) and (3). Hence, thtis would be in line with a weaker output effect for part-ETS firms. However, crucially to link this to within firm leakage we need to see an increase in emissions - and other outcomes - in non ETS plants of ETS firms. We examine this in column (5). For emissions we find a small positive effect. However, it is nowhere near statistical significance. For employment we do not find any significant effect either and the point estimate is equally small and moreover negative.

Hence, it is likely that any heterogeneity between firms with both regulated and unregulated plants, and firms that consist only of regulated plants is due to other technological differences rather than within firm leakage.

[12]i.e. if all nearest neighbour control firms were non-ETS plants in non ETS firms and all emission reductions in ETS plants of multi plant firms are due to leakage to non ETS firms, impact coefficients should double compared to a case were control plants are drawn from non-regulated firms.

5.3 Firm Level Results (forthcoming)

5.3.1 Leakage Across Space? Evidence from Customs Data (forthcoming)

6 Firms' response to carbon pricing - a simple model

Our empirical results so far on the effectiveness of the EU ETS to reduce emissions is somewhat encouraging from a policy perspective. However, from both an academic and policy perspective this introduces an additional puzzle. In the wake of the economic crisis, the permit price in the EU ETS has fallen to record lows. Why would we expect that in such an environment regulated firms would make efforts to reduce emissions any more than unregulated firms? In this section we introduce a simple model that provides a solution to this puzzle.

Most stakeholders expect that prices will increase in the future (Martin et al. 2014). If firms are installing equipment today that will last until prices increase and the amount or quality of this equipment affects emissions, then we can expect reductions today even if current prices are low. We formalize this idea in this section by introducing a simple model of a firm. This presents an interesting paradox: expected future price increases can have a positive effect on current emissions; however, if the increase is so high that firms decide to exit in the future, because any fixed equipment will be used for a shorter amount of time, firms may reduce their investments, thereby increasing current emissions.

This phenomenon is akin to the so called Green Paradox. However it derives not from the presence of an exhaustible polluting resource, but from the long term nature of emission saving investments.

6.1 A simple model

Consider a firm that produces using two production factors: energy E which causes emissions and capital K with the following production function

$$Q = E + K^{\rho}$$

Notice that ρ controls how substitutable the two factors are. If ρ is close to 1 they are highly substitutable, if ρ is close to zero they are not very good substitutes.

Suppose that capital lasts for two periods, hence the firm has a two period planning horizon.

In each period the firm is dealing with an (inverse) demand of the form

$$P = a - bQ$$

Consider first the short run profit maximization for a given amount of capital

$$\max_{Q} (a - bQ) Q - W_E E$$

where W_E is the price of energy.

The first order condition becomes

$$a - 2bQ \leq W_E$$

Consequently

$$E = \begin{cases} \frac{a-W_E}{2b} - K^\rho & \text{if } a - 2bK^\rho > W_E \\ 0 & \text{otherwise} \end{cases}$$

Note that if the capital stock high and/or energy prices are very high the firm will stop using (dirty) energy all-together.

The two period function is consequently

$$\Pi 2 (K) = \begin{cases} \left(\frac{a-W_{E1}}{2}\right)^2 0.5 + \frac{1}{1+r} \left(\frac{a-W_{E2}}{2}\right)^2 0.5 - KW_K - \left[\frac{a-W_{E1}}{2b} - K^\rho\right] W_{E1} - \frac{1}{1+r} \left[\frac{a-W_{E2}}{2b} - K^\rho\right] W_{E2} \\ \text{if } a - 2bK^\rho > W_{E2} \text{ and } \text{ if } a - 2bK^\rho > W_{E1} \\ \\ \left(\frac{a-W_{E1}}{2}\right)^2 0.5 + \frac{1}{1+r} (a - bK^\rho) K^\rho - KW_K - \left[\frac{a-W_{E1}}{2b} - K^\rho\right] W_{E1} \\ \text{if } a - 2bK^\rho \leq W_{E2} \text{ and } \text{ if } a - 2bK^{\rho \frac{1-\mu}{\mu} \rho} > W_{E1} \end{cases}$$

Notice that in the case where the firm uses carbon in both periods we have the following FOC for capital:

$$-W_K + \rho K^{\rho-1} \left[W_{E1} + \frac{1}{1+r} W_{E2}\right] = 0$$

Implying

$$K = \left[\frac{W_{E1} + \frac{1}{1+r} W_{E2}}{W_K} \rho\right]^{\frac{1}{1-\rho}}$$

Hence, the higher the carbon price in period 2 the more capital we are going to invest and as a consequence energy consumption goes down even in period 1. Also note that for any W_{E2} , condition $a - 2bK^\rho > W_{E2}$ is condition is always met if ρ is small enough; i.e. if K and E are bad substitutes we always want to have some energy consumption.

Note that

$$E_1 = \frac{a - W_{E1}}{2b} - \left[\frac{W_{E1} + \frac{1}{1+r}W_{E2}}{W_K}\rho\right]^{\frac{\rho}{1-\rho}} \text{ if } a - 2bK^\rho > W_{E1} \text{ and } a - 2bK^\rho > W_{E2}$$

Alternatively, firms could only be producing in period 1.

Profits when only producing in period 1 are

$$\Pi 1\left(K\right) = \left(\frac{a - W_{E1}}{2}\right)^2 0.5 - KW_K - \left[\frac{a - W_{E1}}{2b} - K^\rho\right] W_{E1}$$

The first order condition for capital investment becomes

$$W_K = \rho K1^{\rho-1} W_{E1}$$

Hence

$$K1 = \left(\frac{\rho W_{E1}}{W_K}\right)^{\frac{1}{1-\rho}}$$

Proposition 1. *If future carbon prices increase beyond a certain threshold, demand is inelastic enough and energy and capital are not very substitutable, then pollution in the present will increase.*

Proof. Notice that

$$\Pi 1\left(K1\right) - \Pi 2(K) = \left(K - K1\right) W_K + \left[K1^\rho - K^\rho\right] W_{E1} + \frac{1}{1+r}\left(\frac{a - W_{E2}}{2}\right)^2 0.5 - \frac{1}{1+r}\left[\frac{a - W_{E2}}{2b} - K^\rho\right] W_E$$

A sufficient condition for firms only producing in period 1 is

$$\Pi 1\left(K\right) - \Pi 2(K) > 0$$

i.e. if one period profits at K are higher than two period profits they must also be higher at $K1$, the one period optimal capital stock. Note that

$$\Pi 1\left(K\right) - \Pi 2(K) = \frac{1}{1+r}\left(\frac{a - W_{E2}}{2}\right)^2 0.5 - \frac{1}{1+r}\left[\frac{a - W_{E2}}{2b} - K^\rho\right] W_{E2}$$

A sufficient condition for that is

$$\frac{1}{1+r}\left(\frac{a - W_{E2}}{2}\right)^2 0.5 - \frac{1}{1+r}\left[\frac{a - W_{E2}}{2b}\right] W_{E2} > 0$$

$$\Leftarrow \left(\frac{a - W_{E2}}{2}\right)^2 0.5 - \left[\frac{a - W_{E2}}{2b}\right] W_{E2} > 0$$

$$\Leftarrow a - \frac{1+b}{b} W_{E2} > 0$$

i.e. $W_{E2} < \frac{ab}{1+b}$ which is true if b is sufficiently small (i.e. demand sufficiently inelastic)

Energy consumption when only producing in period 1 becomes

$$E1_1 \left(W_{E1}\right) = \frac{a - W_{E1}}{2b} - K1^{\rho} = \frac{a - W_{E1}}{2b} - \left(\frac{\rho W_{E1}}{W_K}\right)^{\frac{\rho}{1-\rho}}$$

Let's now compare the energy consumption E of a firm that is exposed to a positive but constant energy price of W_E that of a firm that has to pay a second period price $W_{E2} > W_E$

$$E = \frac{a - W_E}{2b} - \left[\frac{W_E \left(1 + \frac{1}{1+r}\right)}{W_K} \rho\right]^{\frac{\rho}{1-\rho}}$$

Now if we increase consumption to the point where the company goes out of business in period 2:

$$E1_1 \left(W_E\right) - E = -\left(\frac{\rho W_E}{W_K}\right)^{\frac{\rho}{1-\rho}} + \left[\frac{W_E \left(1 + \frac{1}{1+r}\right)}{W_K} \rho\right]^{\frac{\rho}{1-\rho}} = \left(\frac{\rho W_E}{W_K}\right)^{\frac{\rho}{1-\rho}} \left[\left(1 + \frac{1}{1+r}\right)^{\frac{\rho}{1-\rho}} - 1\right]$$

Notice that this is always positive as $\rho < 1$. □

In other words: if the firm does not produce in the second period, then it will invest less capital already in the first period. Hence, firms substitute energy for capital which can increase overall consumption if the demand is sufficiently in-elastic. This ensures that the that the negative output effect because of higher costs does not outweigh the substitution effect.

Figure 1 illustrates this for a particular set of parameter assumptions in a diagram showing different values for energy prices in period 2 on the x-axis and the implied energy consumption in period 1 compared to the consumption if prices were equal in both periods on the y-axis. Hence, we see that as the price increase consumption initially falls by up to 2% Once price increase by about 9 times the period 1 price in period 2 it is no longer profitable to produce in period 2. Hence, energy consumption increases 1% above the baseline level. Notice that the second vertical line in Figure 1 (at around 18) indicates the level at which the firm would stop using energy (provided it would not stop producing in period 2).

Figure 1: The effect of future price increases on emissions today

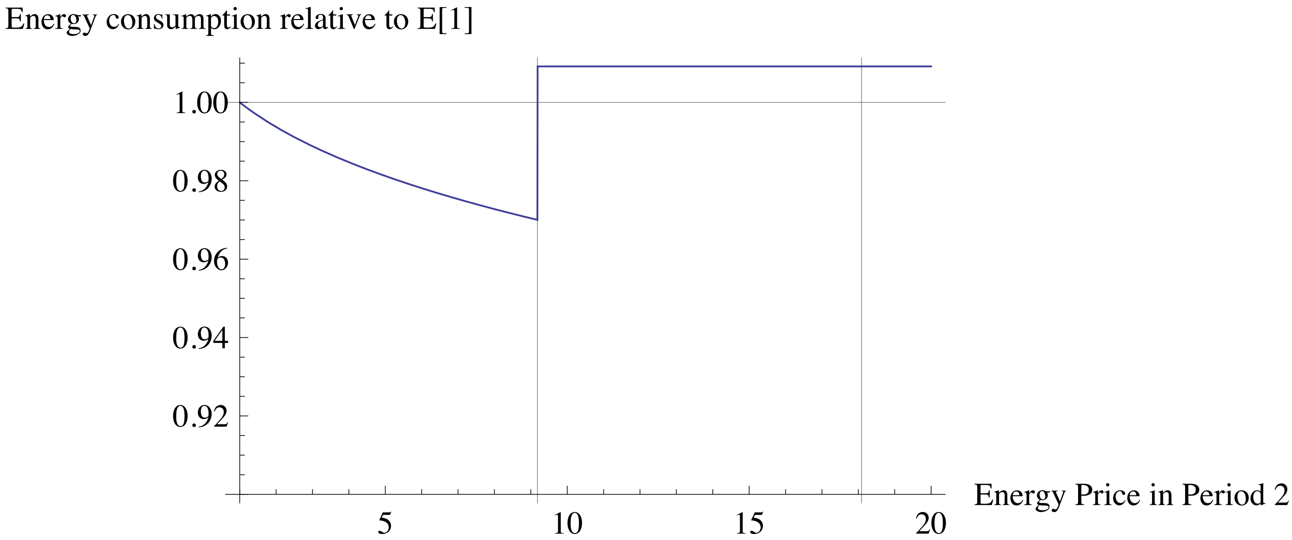

Notes: Parameter assumptions are $\rho = 0.2$, $W_{E1} = 1$, $a = 20$, $b = 0.7$, $W_K = 1$, $r = 10\%$.

6.2 Some implications

Our simple model has a number of implications which we can potentially translate into testable hypothesis. Firstly, it implies that it would be very useful to look at investment and capital as outcome measure as well. This can also help to distinguish if negative employment effects are due to leakage or complementarity with energy. In the latter case we should that ETS causes an increase in investment. In the former case we might expect a reduction in investment - and perhaps an increase in intermediates.

Secondly, it will be interesting to explore if for certain types of firms and sectors we find a positive effect on emissions. We expect this if substitutability is low and demand is inelastic.

A third implication of the model is that firms with a younger capital stock are less likely to respond to changed future price expectations.

The focus of our current research efforts is to test these implications and in doing so provide a more thorough insight into the response of firms' to environmental regulation.

7 Conclusions

This paper provides plant-level evidence of the effectiveness of the EU ETS on the environmental and economic performance of the manufacturing sector in France. Unlike previous research, we use plant-level data to exploit variation in the selection criterion by which plants are required to join the EU ETS. Plant-level data allows us to precisely estimate the impact of the EU ETS by comparing the performance of plants covered by this policy with similar plants that are not. Compared to studies at the aggregate level, which have to identify effects by projecting a baseline into the future, this allows idiosyncratic shocks affecting the economy as a whole and specific sectors over time to be taken into account. It allows us not only to identify and measure the impact of the EU ETS, but also to disentangle the mechanisms by which plants and firms respond.

Our results indicate that the EU ETS has resulted in a significant reduction in GHG emissions (15.7%) within ETS plants, compared to non-ETS plants. So far, we estimate that one of the main drivers of these emission reductions appears to be through changes in the carbon intensity of fuel. In addition, we examine the degree to which firms that have both ETS and non-ETS plants engage in avoidance behaviour such as within-firm leakage – a low-cost avoidance strategy. We find no evidence of within-firm leakage for firms that have both ETS and non-ETS facilities.

In addition to environmental outcomes, it is of policy and academic interest to understand whether the EU ETS has an effect on economic outcomes. Due to its unilateral nature, there are concerns that the EU ETS has had a detrimental impact on the productivity and economic performance of regulated facilities. We observe that ETS plants face significant reductions in employment (-10.4%) compared to non-ETS plants.

Further work is needed to better understand the impacts of the EU ETS on the productivity and economic performance of plants alongside any general equilibrium effects associated with the EU ETS. We also aim to better understand the potential for carbon leakage between markets by incorporating data that links firms to trading partners within and outside the EU.

References (Incomplete)

Åhman, M., D. Burtraw, J. Kruger, and L. Zetterberg (2007): "A Ten-Year Rule to guide the allocation of EU emission allowances," *Energy Policy*, 35, 1718–1730.

Alberola, E., J. Chevallier, and B. Chèze (2009): "Emissions Compliances and Carbon Prices under the EU ETS: A Country Specific Analysis of Industrial Sectors," *Journal of Policy Modeling*, 31, 446–462.

Anger, N. and U. Oberndorfer (2008): "Firm performance and employment in the EU emissions trading scheme: An empirical assessment for Germany," *Energy Policy*, 36, 12–22.

Bushnell, J. B., H. Chong, and E. T. Mansur (2013): "Profiting from Regulation: An Event Study of the EU Carbon Market," *American Economic Journal: Economic Policy*, 5.

Delarue, E. and W. D. D'haeseleer (2007): "Price determination of ETS allowances through the switching level of coal and gas in the power sector," *International Journal of Energy Research*, 31, 1001–1015.

Demailly, D. and P. Quirion (2008): "European Emission Trading Scheme and competitiveness: A case study on the iron and steel industry," *Energy Economics*, 30, 2009–2027.

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, A. D., B. K. Buchner, and C. Carraro, eds. (2007): *Allocation in the European Emissions Trading Scheme: Rights, Rents and Fairness*, Cambridge: Cambridge University Press.

Ellerman, A. D., F. J. Convery, and C. de Perthuis (2010): *Pricing Carbon: The European Union Emissions Trading Scheme*, Cambridge University Press, 1 ed.

Ellerman, A. D. and P. L. Joskow (2008): "The European Union's Emissions Trading System in Perspective," Tech. rep., Pew Center on Global Climate Change, Washington, DC.

Grubb, M., T. L. Brewer, M. Sato, R. Heilmayr, and D. Fazekas (2009): "Climate Policy and Industrial Competitiveness: Ten Insights from Europe on the EU Emissions Trading System," Climate & Energy Paper Series, The German Marshall Fund of the United States.

Martin, R., L. B. de Preux, and U. J. Wagner (2014a): "The impact of a carbon tax on manufacturing: Evidence from microdata," *Journal of Public Economics*, 117, 1–14.

Martin, R., M. Muûls, L. B. de Preux, and U. J. Wagner (2014b): "Industry Compensation Under Relocation Risk: A Firm-Level Analysis of the EU Emissions Trading Scheme," *American Economic Review*, 104, 1–24.

——— (2014c): "On the Empirical Content of Carbon Leakage Criteria in the EU Emissions Trading Scheme." *Ecological Economics*, 105, 78–88.

Martin, R., M. Muûls, and U. Wagner (2013a): "Carbon markets, carbon prices and innovation: Evidence from Interviews with Managers," Paper presented at the Annual Meetings of the American Economic Association, San Diego, January 6.

Martin, R., M. Muûls, and U. J. Wagner (2013b): "The Impact of the EU ETS on Regulated Firms: What is the Evidence after Eight Years?" Available at SSRN: http://ssrn.com/abstract=2344376 or http://dx.doi.org/10.2139/ssrn.2344376.

McKinsey and Ecofys (2006): "Review of Emissions Trading Scheme: Survey highlights," Tech. rep., European Commission Directorate General for Environment, Brussels.

Paltsev, S. V. (2001): "The Kyoto Protocol: Regional and Sectoral Contributions to the Carbon Leakage," *The Energy Journal*, 22, 53–79.

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.

Reinaud, J. (2005): "Industrial competitiveness under the European Union Trading Scheme," Tech. rep., International Energy Agency, Paris.

Trotignon, R. and A. Delbosc (2008): "Allowance trading patterns during the EU ETS trial period: what does the CITL reveal?" .

A Descriptive Statistics

Table 1: Descriptive Statistics - All Plants

	Mean	Std. Dev. (Within)	Std. Dev. (Between)
ETS (Treatment Dummy)	0.051	0.000	0.175
GHG Emissions (1000 tonnes)	10.043	13.445	133.690
Number of Employees	224	77	379
Carbon Factor Intensity	0.050	0.047	0.203
Coal Share	0.016	0.023	0.083
Oil Share	0.183	0.090	0.254
Gas Share	0.332	0.092	0.295
Steam Share	0.005	0.017	0.049
Multiple Plants (Share)	0.356	0.000	0.466
# Observations		102459	
# Plants		11575	
# Firms		9494	

Table 2: Descriptive Statistics - Difference in Means

	Mean ETS	Mean Non-ETS	Difference (Treatment – Control)
GHG Emissions	78.115	6.375	71.739***
Employment	507.735	208.965	298.770***
Carbon Factor Intensity	0.373	0.033	0.340***
Coal Share	0.116	0.010	0.105***
Oil Share	0.161	0.185	-0.023***
Gas Share	0.454	0.326	0.128***
Steam Share	0.015	0.004	0.010***
Multi-Plant	0.591	0.343	0.247***
# Observations	5238	97221	102459
# Plants	368	11207	11575
# Firms	279	9302	9494

Table 3: Descriptive Statistics - ETS plants

	All ETS plants		Only ETS plants		Part ETS plants		Non-ETS plants in ETS firms	
	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev
GHG Emissions	78.115	28.508	79.224	28.101	76.372	29.144	24.435	12.020
Employment	507.735	183.153	478.383	156.131	553.897	219.049	306.629	103.937
Carbon Factor Intensity	0.373	0.156	0.364	0.163	0.386	0.143	0.098	0.047
Coal Share	0.116	0.060	0.086	0.055	0.163	0.066	0.034	0.030
Oil Share	0.161	0.108	0.191	0.110	0.113	0.103	0.138	0.092
Gas Share	0.454	0.115	0.455	0.114	0.452	0.117	0.428	0.092
Steam Share	0.015	0.036	0.016	0.036	0.014	0.037	0.022	0.022
Multiple Plants (Share)	0.591	0.000	0.331	0.000	1.000	0.000	1.000	0.000
# Observations	5238		3202		2036		2485	
# ETS plants	368		227		141		225	
# ETS firms	279		192		87		87	

Figure 2: The share of plants, non-ETS plants, and ETS plants, (by sector)

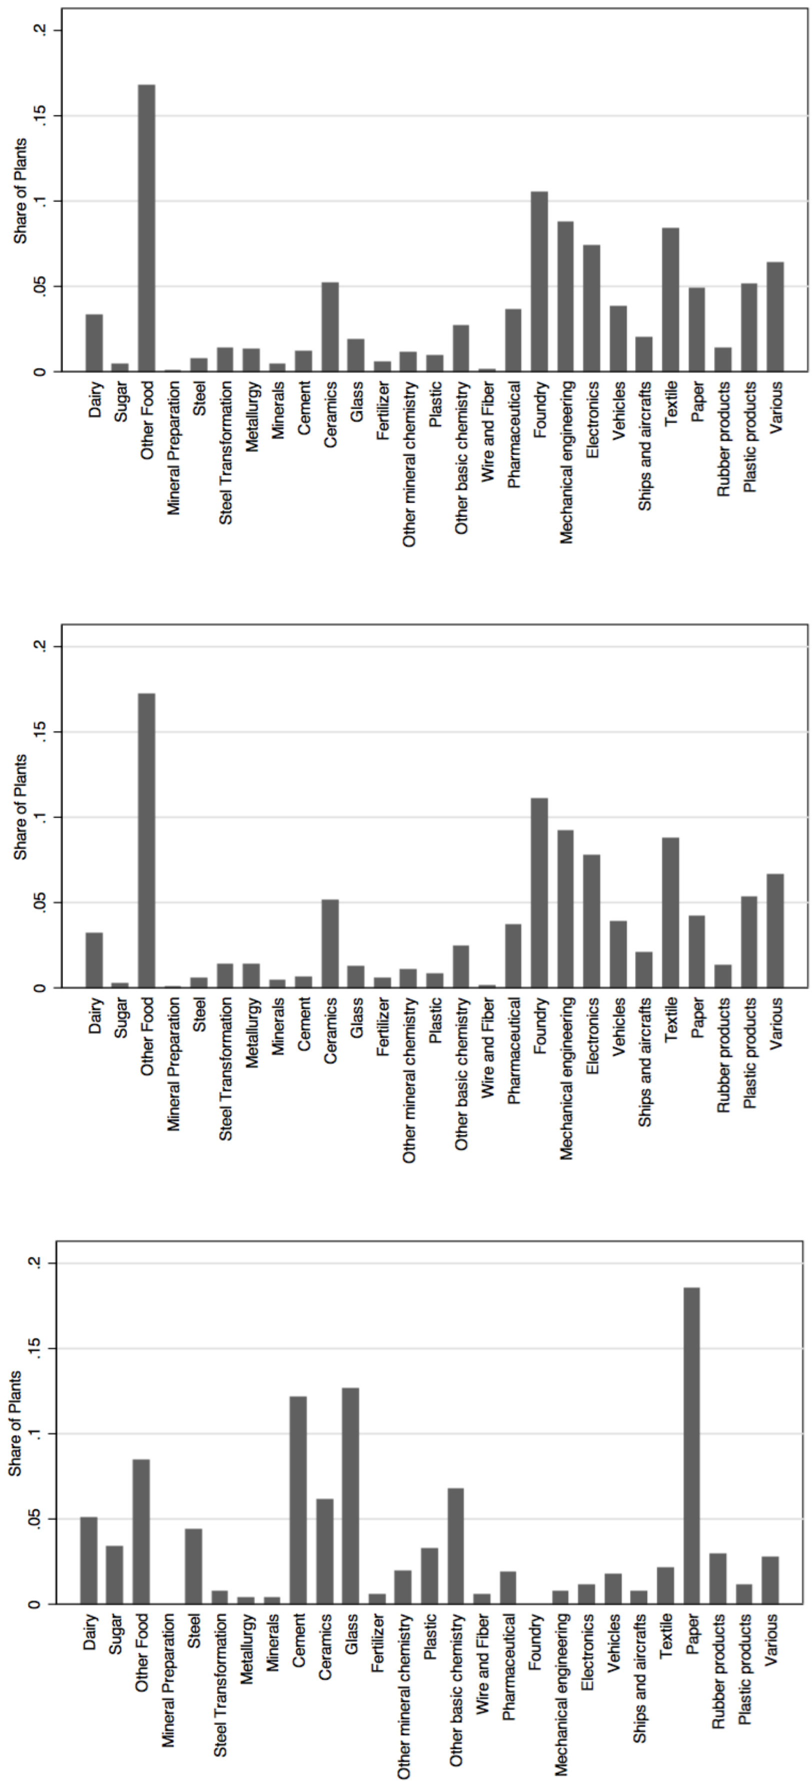

B Results

Table 4: Impact at Plant Level - Results from difference-in-difference propensity score matching estimators (plant level)

	(1) No Restriction	(2) 100%	(3) 50%	(4) 25%	(5) 10%
Panel A: GHG Emissions					
Δln(GHG Emissions)					
Pre-Announcement SATT	0.00388	0.00327	0.00327	0.00305	0.000477
	(0.0147)	(0.0147)	-0.0147	-0.015	-0.0152
Announcement Phase SATT	0.00382	0.00418	0.00402	0.00443	0.00575
	(0.0131)	(0.0132)	(0.0132)	(0.0134)	(0.0141)
Phase I (2005-2007) SATT	-0.0258	-0.0278	-0.0245	-0.0272	-0.0358
	(0.0252)	(0.0255)	(0.0261)	(0.0266)	(0.0282)
Phase II (2008-2010) SATT	-0.133***	-0.139***	-0.149***	-0.146***	-0.157***
	(0.0349)	(0.035)	(0.0351)	(0.0378)	(0.0418)
Panel B: Employment					
Δln(Employment)					
Pre-Announcement SATT	-0.00822	-0.0082	-0.0082	-0.00814	-0.00769
	(0.00993)	(0.01)	(0.01)	(0.0101)	(0.0103)
Announcement Phase SATT	-0.0477***	-0.0457***	-0.0446***	-0.0418***	-0.0440***
	(0.0141)	(0.0141)	(0.0142)	(0.0141)	(0.0147)
Phase I (2005-2007) SATT	-0.0425*	-0.0412*	-0.0369	-0.0422	-0.0491*
	(0.0247)	(0.0249)	(0.0253)	(0.0257)	(0.0273)
Phase II (2008-2010) SATT	-0.0908***	-0.0896***	-0.0904***	-0.105***	-0.104***
	(0.034)	(0.0335)	(0.0332)	(0.0342)	(0.0369)
Panel B: Carbon Intensity					
Δln(GHG/Employment)					
Pre-Announcement SATT	0.0121	0.0115	0.0115	0.0112	0.00816
	(0.016)	(0.0161)	(0.0161)	(0.0163)	(0.0168)
Announcement Phase SATT	0.0515***	0.0499***	0.0486***	0.0462***	0.0498***
	(0.0157)	(0.0157)	(0.0157)	(0.0159)	(0.0166)
Phase I (2005-2007) SATT	0.0167	0.0135	0.0123	0.0149	0.0134
	(0.0271)	(0.0275)	(0.028)	(0.0289)	(0.0309)
Phase II (2008-2010) SATT	-0.0422	-0.0494	-0.0584	-0.0415	-0.0529
	(0.0356)	(0.0359)	(0.0368)	(0.0389)	(0.0441)
Observations	5238	5169	5122	4965	4610
Clusters	279	276	276	272	265

Notes: Treatment plants are matched to control plants based on carbon intensity, with exact matching on sectors at the 2 digit level. Column 1 has no restrictions on the distance between neighbours. Column 2 restricts the distance of the nearest neighbour to 100%. Column 3 restricts the distance of the nearest neighbour to 50%. Column 4 restricts the distance of the nearest neighbour to 25%. Column 5 restricts the distance of the nearest neighbour to 10%. Significance levels are indicated as * 0.10 ** 0.05 *** 0.001. Robust Standard errors, clustered at the firm level, are in parentheses.

Figure 3: Difference in log(GHG) emissions, log(Employment), and log(Carbon Intensity) between treatment and comparison plants over time. Base Year 1999 (Maximum Distance 10%)

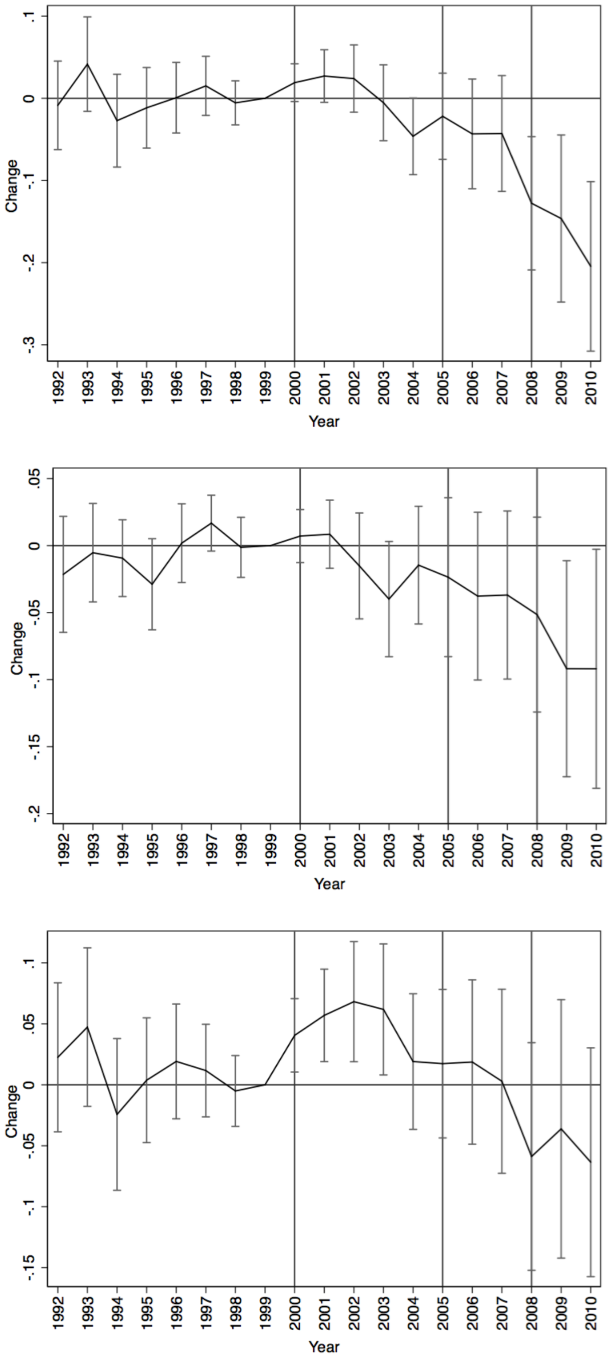

The Composition of Emissions

Table 5: Impact on Fuel Mix - Results from difference-in-difference propensity score matching estimators (plant level)

	(1)	(2)	(3)	(4)
	Coal Share	Oil Share	Gas Share	Steam Share
Pre-Announcement SATT	-0.00167	-0.0062	0.0125	0.000287
	(0.00505)	(0.00813)	(0.00923)	(0.00149)
Announcement Phase SATT	-0.0153**	-0.0109*	0.0207**	0.00555*
	(0.00702)	(0.00609)	(0.00939)	(0.00328)
Phase I (2005-2007) SATT	-0.0275**	-0.0238*	0.0422**	0.00965
	(0.0125)	(0.014)	(0.0194)	(0.00678)
Phase II (2008-2010) SATT	-0.0477***	-0.0492**	0.0766***	0.0112
	(0.0164)	(0.0199)	(0.0258)	(0.00748)
# Observations	4610	4610	4610	4610
# Clusters	265	265	265	265

Notes: Treatment plants are matched to control plants based on carbon intensity, with exact matching on sectors at the 2 digit level. All columns restricts the maximum distance to 10%. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Robust Standard errors, clustered at the firm level, are in parentheses.

The Composition of the Firm

Table 6: Impact by type of firm - Results from difference-in-difference propensity score matching estimators (plant level)

	(1)	(2)	(3)	(4)	(5)
Panel A: GHG Emissions					
∆ln(GHG Emissions)					
Pre-Announcement SATT	0.000477	-0.00302	-0.00399	0.00456	0.0604***
	-0.0152	(0.0302)	(0.0304)	(0.0162)	(0.0226)
Announcement Phase SATT	0.00575	-0.0186	-0.0305	0.0131	0.0181
	-0.0141	(0.0237)	(0.0222)	(0.0183)	(0.0227)
Phase I (2005-2007) SATT	-0.0358	-0.0164	-0.0038	-0.0661*	0.0525
	-0.0282	(0.0403)	(0.0411)	(0.0395)	(0.0507)
Phase II (2008-2010) SATT	-0.157***	-0.111*	-0.0966	-0.210***	0.0185
	-0.0418	(0.0645)	(0.0691)	(0.0552)	(0.0661)
ln(GHG Emissions) 1999	3.481	3.508	3.508	3.463	1.661
Panel B: Employment					
∆ln(Employment)					
Pre-Announcement SATT	-0.00769	0.0153	0.0187	-0.0220*	0.0339*
	-0.0103	(0.0157)	(0.0158)	(0.0132)	(0.0201)
Announcement Phase SATT	-0.0440***	-0.0600**	-0.0773***	-0.0355*	-0.0497**
	-0.0147	(0.0255)	(0.0241)	(0.019)	(0.0233)
Phase I (2005-2007) SATT	-0.0491*	-0.0663*	-0.0609*	-0.0487	-0.0122
	-0.0273	(0.0366)	(0.0365)	(0.0399)	(0.0438)
Phase II (2008-2010) SATT	-0.104***	-0.119**	-0.123**	-0.106**	-0.00922
	-0.0369	(0.052)	(0.0533)	(0.053)	(0.0476)
ln(Employment) 1999	5.570	5.657	5.657	5.514	5.240
Panel B: Carbon Intensity					
∆ln(GHG/Employment)					
Pre-Announcement SATT	0.00816	-0.0183	-0.0227	0.0266*	0.0265
	-0.0168	(0.0354)	(0.0361)	(0.0159)	(0.0224)
Announcement Phase SATT	0.0498***	0.0414	0.0468*	0.0486**	0.0678**
	-0.0166	(0.0263)	(0.0268)	(0.0223)	(0.0279)
Phase I (2005-2007) SATT	0.0134	0.0499	0.0571	-0.0174	0.0646
	-0.0309	(0.0509)	(0.0523)	(0.0433)	(0.0433)
Phase II (2008-2010) SATT	-0.0529	0.00799	0.0262	-0.104*	0.0278
	-0.0441	(0.0732)	(0.0755)	(0.0592)	(0.0646)
ln(Carbon Intensity) 1999	-2.089	-2.149	-2.149	-2.051	-3.578
# Observations	4610	1770	1707	2736	2248
# Clusters	265	83	81	180	84
Treatment Group	ETS plants	Part ETS	Part ETS	Only ETS	Non-ETS in ETS
Control Group	Non-ETS plants	Non-ETS plants	Non-ETS plants	Non-ETS plants	Non-ETS plants
Non-ETS plants in ETS firms included in control group	Yes	Yes	No	N/A	--

Notes: Treatment plants are matched to control plants based on carbon intensity, with exact matching on sectors at the 2 digit level. All columns restricts the maximum distance to 10%. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Robust Standard errors, clustered at the firm level, are in parentheses.

C Balance Tests

The following figures, present balance tests for the matching covariates used. In addition, we present results demonstrating robustness to the restriction of the maximum distance between the treatment observation and its nearest neighbor. The base year used in all balance tests is 2000.

Figure 4: Distribution of pre-treatment log(Carbon Intensity) for ETS plants and non-ETS plants: before matching and after matching

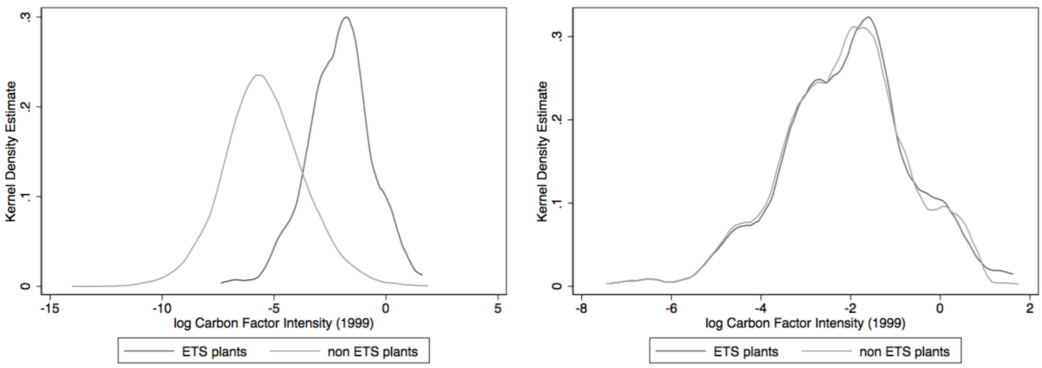

Figure 5: QQ plot of pre-treatment log(Carbon Intensity) for ETS plants against pre-treatment log(Carbon Intensity) for non-ETS firms before matching and after matching

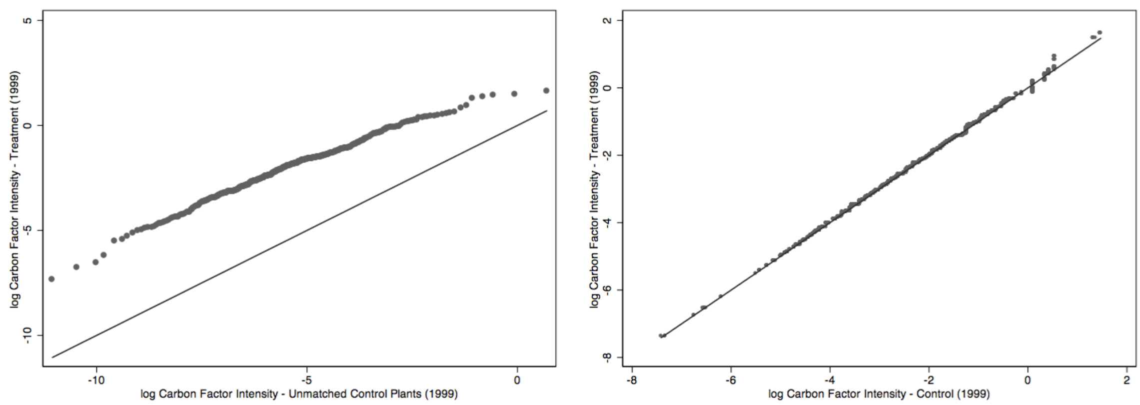

By comparing the plots in Figure 4 we can see the improvements that matching has made in the distributions of the treatment and control groups. The QQ plots in Figure 5 further demonstrate the improvements that matching can make. Panel A in Figure 5 shows the probability distributions of the treatment and control groups without matching; panel B the distributions when matching on carbon intensity (50% maximum distance); panel C the distributions when matching on carbon intensity (10% maximum distance).

The QQ plots in panel B of Figure 5 indicates that there may be some skewness in the distribution to the right. In spite of this the balance is an improvement upon the unconditional difference.

In the graphs below, we visualize the change in the sample used when we reduce the maximum distance between treatment units and their nearest neighbor at intervals of 100%, 50%, and 10%.

Figure 6: Nearest neighbours included when the maximum distance between treatment and control is restricted to 10%