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

Ulrich Wagner* Mirabelle Muûls† Ralf Martin‡ Jonathan Colmer§

1st April 2014

Abstract

Using comprehensive firm and plant level data for more than 4,500 French manufacturing firms, this paper examines the impact of the European Union Emissions Trading Scheme. The EU ETS is the EUs flagship climate policy, establishing since 2005 a carbon market across Europe with more than 11,400 plants in 31 countries regulated at present. Our results suggest that ETS regulated manufacturing plants in France reduced emissions by an average of 15-20%, a significant amount. 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 from a climate policy point of view, we also find a significant reduction of employment in regulated firms of about 7% in Phase II.

*Universidad Carlos III Madrid, Department of Economics, Calle Madrid, 126 28903 Getafe, Spain and Research Associate, Centre for Economic Performance, London School of Economics and Political Science, phone: +34 (0) 916 24 8488, email: uwagner@eco.uc3m.es

†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.

‡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.

§Grantham Research Institute on Climate Change and the Environment, London School of Economics and Political Science, Houghton Street, London WC2A 2AE. E-mail: j.m.colmer@lse.ac.uk. Support from the ESRC Centre for Climate Change Economics and Policy and the Grantham Foundation is gratefully acknowledged.

1 Introduction

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 the Emissions Trading System (EU ETS), the first international carbon trading scheme in the world. 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 EEA states, the EU ETS covers large plants from CO_2-emission intensive industrial sectors, namely power generation, mineral oil refineries, coke ovens, iron and steel and factories producing cement, glass, lime, brick, ceramics, pulp and paper, and all combustion activities with a rated thermal input exceeding above 20MWh. 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 is paper estimates the causal impact of the EU ETS using comprehensive firm and plant level data. Plant level data allows us to compare the performance of plants covered by the EU ETS 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.

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.

Because of the unilateral nature of the EU ETS, there is also considerable concern regarding its impact on economic outcomes, specifically impacts on the competitiveness, economic performance and employment of regulated companies. 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 EU ETS-regulated manufacturing plants reduced emissions by an average of 15-20%, a significant amount. 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 from a climate policy point of view, we also find a significant reduction of employment in regulated firms of about 7% in Phase II.. Because these results are relative to non-regulated firms we have to be cautious in making claims about aggregate effects. Emission reductions at the national or global level could be lower than 20% if there are leakage effects. These could arise either as leakage from regulated EU firms to non-regulated sectors within the EU or outside the EU. Analysis of further outcome variables in future research could shed further light on this issue as we discuss in detail below. However, it is permissible to consider our results as an upper bound on the effects of the EU ETS on emissions. Arising from concerns surrounding leakage, we aim to disentangle the channels through which these reductions take place. We find little evidence of within-firm leakage, presumably the least cost leakage strategy available to firms that have both ETS and non-ETS facilities, and observe that reductions arise predominantly through adjustments relating to the composition of emissions i.e. fuel mix. Given the lack of leakage within-firm, these results could mitigate concerns of strategic behaviour and its implications for leakage between markets. However, additional opportunities in markets outside of the ETS such as cheaper energy costs, may still raise concerns.

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 and the results on both energy consumption and efficiency and section 5 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 & 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 et al., 2013c).

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 see McKinsey & 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 et al. (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. (2013d) review this literature in detail. Demailly & Quirion

(2008) and Anger & 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 and 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 and at the same time a lower risk of carbon leakage and loss of competitiveness (Martin et al., 2013b and Martin et al., 2013a). Apart from that, researchers have been keen to evaluate the effectiveness of the EU ETS at incentivizing carbon abatement. Ellerman & Buchner (2008) and work by Delarue & 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 & Joskow, 2008). [1]

Wagner et al., 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 employment was not affected by the EU ETS, they find moderate but statistically significant increases in the value of output and exports during the second phase. In contrast to their firm-level analysis, this paper also analyses the policy impact at the sub-firm level.

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 in terms of discerning the causal effect of the EU ETS from other concurrent events. Below, we will address this problem by making use of longitudinal observations at the level of the individual business enterprise, for both ETS and non-ETS firms to study the effects of the policy on energy use and economic performance.

3 Data

3.1 L'Enquete Annuelle sur les Consommations d'Energie dans l'Industrie (EACEI)

The EACEI (Annual survey of energy consumptions in the industry) is a survey conducted

[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 et al., 2011).

until 2008 by the Service des études et des statistiques industrielles (SESSI)[2] of the ministry in charge of manufacturing.[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 from the chemical decomposition 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 of 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 of this sample 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 L'Enquete Annuelle des Entreprise (EAE)

The Enquete Annuelle des Entreprise, French annual business surveys data, also conducted by SESSI, are available for the period 1993 to 2007 . Firm and plant level data is available for firms with more than 20 employees and all the plants of those firms. At the firm level, the dataset provides balance-sheet data including 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 but includes main variables of interest.

3.3 The Community Independent 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

[2]Since 2009 the EACEI is conducted by the Statistical direction of INSEE, the National Institute of Statistics and Economic Studies

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

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

registries interlink with the Community Independent Transaction Log (or CITL), operated by the Commission, which will record and check every transaction. It is now centralised as the European Union Transaction Log (or EUTL). Each of the 12,000 plants has an "operator holding account" in its national registry, into which its own allowances are being 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" as well as the compliance position of each plant. Records of other types of transactions will be 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 5,957 plants within 4,589 firms. 287 firms and 384 plants 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 are: within firms with multiple plants that only have ETS plants; firms with both ETS and non-ETS plants; single plant ETS firms.

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 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

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

sample average treatment effect on the treated (SATT):

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

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 of 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 during 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 number of strategies that condition on observable variables.

4.1 Parametric conditioning strategies .

Within a parametric framework, we employ two strategies to reduce any bias introduced from differences in the characteristics between ETS and non-ETS plants. First we use ordinary least squares (OLS) to control for observable factors that affect plant-level emissions and energy use. Secondly, we impose a common support on the sample of the treatment and control groups and control for observables in order to provide a similar distribution of covariates between ETS and non-ETS plants.

Under both conditions we use OLS to estimate the following specification:

$$\Delta Y_i = \alpha \cdot ETS_i + \beta X_i' + \epsilon_i,$$

where $\Delta Y_i = Y_{it_1} - Y_{it_0}$, the difference in the outcome variable between the post-treatment and pre-treatment time period, X_i, a vector of observable covariate, including pre-treatment plant-specific economic measures. Finally, α captures the average treatment effect on the treated of

[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.

the EU ETS on changes in Y_i over time, and ϵ_i the error term is assumed to be independent of the treatment indicator, ETS_i, and covariates in X_i.

However, there are a number of concerns which may accentuate the bias in our results, even after controlling for observable differences between treated and untreated plants. These concerns may be warranted if there are differences in the distributions of the vector of control variables, X_i, between ETS and non-ETS plants. If the control variables for non-ETS plants are not appropriately reweighted to impose a common support to control for differences in the distribution, this will bias the α_{ATT}.

As a first effort to mitigate the effects of biases, resulting from distributional concerns we first restrict the sample to provide a measure of common support between ETS and non-ETS plants. As an alternative, and more robust method of mitigating these concerns we implement semi-parametric matching estimators in the spirit of Heckman, Ichimura, and Todd (1997; 1998), Blundell et al. (2001) and Abadie (2005).

4.2 Semi-parametric conditioning strategies.

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:

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

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}. Weights for control plants are constructed using a process of nearest neighbour matching, where similarity is based on X.

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 trends assumption is satisfied, other assumptions, such as SUTVA are not testable. However, we can implement indirect tests to evaluate whether these assumptions are plausible. This is of importance as it might be the case that within a firm, emissions are reallocated from ETS plants to non-ETS plants. By, aggregating our analysis to the firm level we can check for potential spillovers within firm, although we are unable to rule out spillovers between firms. 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 version of the paper we report results from the semi-parametric difference-in-difference approach. For each EU ETS plant × year we identify the closest match based on a propensity score. The propensity score is estimated using the carbon intensity of each plant in the year 2000, the announcement year of the EU ETS. We also match exactly on the 2-digit sector of each plant.[7] Balance tests are reported in the appendix.

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 2000 are defined as the pre-announcement phase, a test of the parallel trends assumption. Between 2000 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 2005-2007 and Phase II is defined as 2008-2010[8]. The coefficient in table 4 reports the average of each of these Phases, the average treatment effect on the treated. In addition, the average treatment effect on the treated for each year is reported graphically after each table.

5.1 Plant Level Results

Table 4 presents the main results for this paper, applying a 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 relative to the announcement of the EU ETS in 2000 for a different outcome variable: GHG emissions, Employment, and Carbon Intensity.

In panel A, we observe the effect of the EU ETS on the GHG emissions of ETS plants, relative to non-ETS plants. Of importance we note that during the pre-announcement period there appears to be no significant differences 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, relative to the baseline year 2000, however, the coefficients are negative and increasing over time. 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 13.5-19.8%, with the effect size increasing as we restrict the maximum distance between matching partners, reducing measurement error, and consequently attenuation bias in the reported coefficients.

Panel B, reports the effects of the EU ETS on employment within ETS plants, compared to

[7] In the event that plants report different sectors over time, we use the sector that is reported most often.
[8] Phase II carried on until 2012, but the confidential data for French firms is only available until 2010

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 the announcement of the EU ETS reduced employment by around 2%, while Phase II of the EU ETS reduced employment by around 7%. However, because these results are relative to non-regulated plants we have to be cautious in making claims about aggregate effects. With regards to the employment effects and their welfare consequences, it is unclear whether the EU ETS increased unemployment, or whether employees can easily gain work locally in non-ETS plants. We find no effects on employment in Phase I. Furthermore, the lack of an effect in the pre-announcement phase is consistent with the assumption of parallel trends.

The EU ETS appears to have had no effect during the Phase I on neither 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 free allocation of permits which resulted in a plummeting of the price on the market after a few years, thus reducing the incentive to make investments to reduce emissions. In addition, firms may have questioned the extension of the EU ETS into Phase II.

Unfortunately, due to data limitations, we are unable to examine the effect of the EU ETS on carbon intensity, 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 reduced the carbon intensity of production by 8-12% during Phase II with no robust effects during the other phases.[9] This indicates that the reductions in emissions are significantly greater than the reductions in employment, improving the efficiency of production with respect to GHG emissions.

5.2 Firm Level Results

Table 5 aggregates the plant level data to each firm in order to test the robustness of these results and identify the potential within-firm leakage or reallocation.In panel A, we observe similar results to those observed at the plant level. There are no significant differences in the emissions between ETS and non-ETS facilities prior to the second phase. In Phase II we observe significant reductions in emissions, however, the magnitude of this effect is markedly decreased to around a 10% reduction in emissions. It is unclear whether this is the result of within-firm leakage, or the result of measurement error attenuating the magnitude of the effect. Measurement error is introduced when aggregating to the firm level as the selection criterion for EU ETS membership is at the plant level. Consequently, if firms contain both non-ETS and ETS facilities we are assigning the non-ETS facilities ETS status when aggregating to the firm level. However, it may still be the case that firms reallocate emissions within-firm as this would be the least cost strategy of mitigating the effects of the EU ETS. This provides an indirect test of the potential for leakage between markets. With rational firms, evidence of

[9] We observe a weakly significant increase in carbon intensity during the announcement phase, but this is not robust once we restrict the maximum distance between nearest neighbours.

leakage within firms would point to the possibility of leakage between markets. However, if there is no evidence of within-firm carbon leakage , presumably the least-cost option, then an argument could be made that firms are unlikely to reallocate emissions overseas. The following section, aims to decompose the effect of the EU ETS on emissions through an examination of firm structure as well as the composition of emissions in terms of fuel share.

In contrast to the plant level results we find no effect of the EU ETS during Phase II on firm-level employment. As a result, we also observe no effect on carbon intensity. As discussed, measurement error is introduced as a result of aggregating to the firm level so it may well be the case that the effects on labour are too noisy at the firm level to identify, supported by the imprecision of the coefficients, which are of similar size to the plant level coefficients. However, it may well be possible that firms re-allocate labour within the firm, a reasonable outcome even in the short-run. As with the emissions results, we aim to better understand whether within-firm reallocation is possible through an examination of firm structure.

5.3 Disentangling the Channels

5.3.1 The Composition of the Firm

Table 6 presents the results of our analysis to disentangle the plant level results, exploiting the structure of firms, to uncover evidence of strategic behavior and within-firm leakage.[10] In column (1) and (2), we report the impact of the EU ETS on ETS plants that are part of firms that contain non-ETS plants relative to all non-ETS plants. We observe a smaller coefficient size equal to around 11% in Phase II, relative to the effect across all plants, however this effect is insignificant, most likely due to a large reduction in power associated with sample size. This may capture the opportunities that these multi-plant ETS firms have, relative to firms that only have ETS plants, in reallocating emissions within-firm. To test this hypothesis further columns (3) and (4) presents the impact of the EU ETS on ETS plants that are part of firms that contain non-ETS plants relative to non-ETS plants that are not part of ETS firms. This removes the opportunity for ETS plants to be matched with other plants within the same firm. If there is evidence of carbon leakage within-firms we would expect that removing this option would reduce the size of the coefficient as the difference would only capture reductions from ETS firms, as opposed to increases by control plants. However, we observe that there are is effectively no difference in the coefficients between columns (1) and (3) and columns (2) and (4). This appears to indicate that in columns (1) and (2) the treatment plants are not being matched to other plants within the firm. While this test doesn't support within-firm leakage, the results for these plants are insignificant and roughly 50% smaller than the average effect across plants. Consequently, further testing is required to confirm the absence of within-firm leakage. Columns (5) and (6) report the results for firms that have no opportunity for reallocating emissions within-firm i.e. single plant ETS firms and multi-plant ETS firms with

[10] The results for each firm structure are reported twice. The first column restricts the maximum distance between neighbours to 50%. The second column restricts the maximum distance between neighbours to 10%.

only ETS plants. If within-firm leakage occurs within multi-plant firms that have both ETS and non-ETS plants, we should expect the effect size for firms that only have ETS plants to be much larger. We observe that during Phase II ETS plants reduce emissions by between 12-16%, relative to non-ETS plants. The magnitude of the coefficients in columns (5) and (6) are slightly larger than the magnitude of the coefficients in columns (1-4), indicating that if there is within-firm leakage it is equivalent to an effect size equal to 2-5%.

A sufficient condition for the presence of within-firm leakage is to test the impact of the EU ETS on non-ETS plants in ETS firms, reported in columns (7) and (8). If there is evidence of within-firm leakage we should observe a significant increase in emissions for these plants. In both column (7) and (8) we observe no significant effect of the EU ETS on emissions. Furthermore, as we restrict the distance between neighbors from 50% to 10%, the pre-treatment trends improve and the magnitude of the coefficients decrease from a 3% increase in emissions to a 0.8% increase in emissions, indicating the precision of the zero effect. Combined, these results do not provide convincing evidence of within-firm reallocation of emissions, providing supporting evidence for an argument against between-market carbon-leakage. However, the sample size in columns (5) and (6) are roughly 3 times larger than columns (1-4) and columns (7) and (8) providing much greater power to identify a statistically significant effect. If with an increase in sample size the effects observed in columns (7) and (8) were significant, the size of within-firm leakage appears to be very small, around 5% of emissions for multi-plant firms that have non-ETS plants. From the descriptive statistics in Table 3, these plants account for 30% of ETS plants. Consequently, with a leakage effect of 5% within these plants, this corresponds to a leakage effect equal to 1.65% across the EU ETS.

The evidence appears to suggest that there is little or no leakage within-firms, indicating that between-market leakage may also be limited. However, it is important to note that there may be additional benefits to reallocating emissions to unregulated markets, such as reductions in energy costs. Furthermore, we are unable to directly reject the hypothesis that firms with no within-firm reallocation opportunities reallocate emissions between markets. It is important to note that it is still possible, even without within-firm carbon leakage, for between-market carbon leakage to occur from single-plant and multi-plant firms that have only ETS plants.

5.3.2 The Composition of Emissions

In order to better understand whether reductions in emissions are real reductions or the result of leakage, we explore the composition of emissions. This allows us examine the degree to which emissions are reduced through changes in the fuel mix. Table 7 reports the impact of the EU ETS on the role that coal, oil, gas and steam as a share of each plant's fuel mix. The result that stands out most is a reduction in coal by ETS plants during Phase II. We observe that ETS plants, relative to non-ETS plants reduce their share of coal by around 2%. We see similar reductions in the share of oil, and increases in the share of gas, and small increases in steam use, however, these are not statistically significant. This indicates that the carbon intensity of the fuel mix in ETS plants falls during Phase II of the EU ETS.

6 Conclusions

This paper provides plant-level evidence of the effectiveness of the EU ETS on environmental and economics outcomes. 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. This not only allows us to identify and measure the impact of the EU ETS, but also 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 within ETS plants, close to 20%, relative to non-ETS plants. We observe that the driver of these emission reductions appears to be through changes in the carbon intensity of fuel used by the plants. In addition, we examine the degree to which firms that have both ETS and non-ETS plants engage in strategic behaviour resulting in within-firm leakage, the least-cost mitigation strategy. We find little evidence to support this hypothesis, finding that, at most, within-firm carbon leakage would account for 5 percentage points of the emission reductions for multi-plant firms that have non-ETS plants. However, this result is statistically insignificant and accounts for only a 1-2% reduction in emissions, relative to non-ETS plants, 10-20% of the overall emission reductions. Given, the lack of evidence for within-firm leakage, this begs the question as to how much between-market leakage occurs. Future work, aims to examine the effect of the EU ETS on imports and exports within- and outside of the European Union linking firms with their trading partners.

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 competitiveness, economic performance, and employment of regulated companies. We observe that ETS plants face significant reductions in employment, equal to around 7%, relative to non-ETS facilities.

Our results indicate that in spite of concerns regarding the effectiveness of the EU ETS, there are considerable reductions in emissions. While there appear to be some reductions in employment, further work is needed to better understand these impacts on competitiveness and economic performance alongside any general equilibrium effects associated with the EU ETS. Furthermore, a better understand the potential for carbon leakage using import and export data linking firms to trading partners within and outside of the EU.

References

Åhman, M., Burtraw, D., Kruger, J., & Zetterberg, L. (2007). A ten-year rule to guide the allocation of eu emission allowances. *Energy Policy*, 35(3), 1718–1730.

Alberola, E., Chevallier, J., & Chèze, B. (2009). Emissions compliances and carbon prices under the eu ets: A country specific analysis of industrial sectors. *Journal of Policy Modeling*, 31(3), 446–462.

Anger, N. & Oberndorfer, U. (2008). Firm performance and employment in the eu emissions trading scheme: An empirical assessment for germany. *Energy Policy*, 36(1), 12–22.

Delarue, E. & D'haeseleer, W. D. (2007). Price determination of ets allowances through the switching level of coal and gas in the power sector. *International Journal of Energy Research*, 31(11), 1001–1015.

Demailly, D. & Quirion, P. (2008). European emission trading scheme and competitiveness: A case study on the iron and steel industry. *Energy Economics*, 30(4), 2009–2027.

Ellerman, A. D. & Buchner, B. K. (2008). Over-allocation or abatement? a preliminary analysis of the eu ets based on the 2005/06 emissions data. *Environmental and Resource Economics*, 41(2), 267–287.

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

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

Ellerman, A. D. & Joskow, P. L. (2008). *The European Union's Emissions Trading System in Perspective*. Technical report, Pew Center on Global Climate Change, Washington, DC.

Grubb, M., Brewer, T. L., Sato, M., Heilmayr, R., & Fazekas, D. (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., de Preux, L. B., & Wagner, U. J. (2011). *The Impacts of the Climate Change Levy on Manufacturing: Evidence from Microdata*. Working Paper 17446, National Bureau of Economic Research.

Martin, R., Muûls, M., de Preux, L., & Wagner, U. J. (2013a). On the Empirical Content of Carbon Leakage Criteria in the EU Emissions Trading Scheme. Available at SSRN: http://ssrn.com/abstract=2362820 or http://dx.doi.org/10.2139/ssrn.2362820.

Martin, R., Muûls, M., de Preux, L. B., & Wagner, U. J. (2013b). Industry compensation under relocation risk: A firm-level analysis of the EU Emissions Trading Scheme. *American Economic Review*, forthcoming.

Martin, R., Muûls, M., & Wagner, U. (2013c). 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., Muûls, M., & Wagner, U. J. (2013d). 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 & Ecofys (2006). *Review of Emissions Trading Scheme: Survey highlights*. Technical report, 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(4), 53–79.

Reinaud, J. (2005). *Industrial competitiveness under the European Union Trading Scheme*. Technical report, International Energy Agency, Paris.

Trotignon, R. & Delbosc, A. (2008). Allowance trading patterns during the eu ets trial period: what does the citl reveal?

Wagner, U. J., Rehdanz, K., & Petrick, S. (2014). The impact of carbon trading on industry: Evidence from German manufacturing firms. Mimeograph.

Descriptive Statistics

Table 1: Descriptive statistics - All Plants

	Mean	Std. Dev	Observations	No. of Plants	No. of Firms
All plants					
GHG Emissions	14.462	151.117	66,596	5,957	4,589
Employment	306.373	508.492	66,596	5,957	4,589
Multiplant Firm	0.441	0.496	66,596	5,957	4,589
Carbon Intensity	0.066	0.298	66,596	5,957	4,589
Coal Share	0.007	0.067	66,596	5,957	4,589
Oil Share	0.143	0.232	66,596	5,957	4,589
Gas Share	0.363	0.294	66,596	5,957	4,589
Steam Share	0.007	0.062	66,596	5,957	4,589
Multi-plant firms					
GHG Emissions	20.231	215.524	29,385	2,647	1,279
Employment	369.701	610.847	29,385	2,647	1,279
Carbon Intensity	0.075	0.259	29,385	2,647	1,279
Coal Share	0.009	0.074	29,385	2,647	1,279
Oil Share	0.135	0.229	29,385	2,647	1,279
Gas Share	0.368	0.292	29,385	2,647	1,279
Steam Share	0.011	0.076	29,385	2,647	1,279
Single-plant firms					
GHG Emissions	9.906	64.364	37,211	3,310	3,310
Employment	256.364	403.026	37,211	3,310	3,310
Carbon Intensity	0.059	0.326	37,211	3,310	3,310
Coal Share	0.005	0.060	37,211	3,310	3,310
Oil Share	0.149	0.235	37,211	3,310	3,310
Gas Share	0.358	0.296	37,211	3,310	3,310
Steam Share	0.004	0.048	37,211	3,310	3,310

Notes: Descriptive statistics are for all available years.

Table 2: Descriptive statistics - Difference in Means

	Mean ETS	Mean non-ETS	Difference (ETS - non-ETS)
All plants			
GHG Emissions	104.341	6.568	97.773***
	(6.861)	(0.173)	(2.115)
Employment	554.751	284.557	270.193***
	(18.391)	(1.374)	(7.156)
Multiplant Firm	0.608	0.426	0.181***
	(0.006)	(0.001)	(0.007)
Carbon Intensity	0.377	0.039	0.337***
	(0.009)	(0.0008)	(0.004)
Coal Share	0.038	0.004	0.034***
	(0.001)	(0.0002)	(0.0009)
Oil Share	0.157	0.142	0.014***
	(0.003)	(0.0009)	(0.003)
Gas Share	0.452	0.355	0.097***
	(0.001)	(0.004)	(0.004)
Steam Share	0.015	0.006	0.009***
	(0.001)	(0.0002)	(0.0008)
Observations	5,377	61,219	66,596
Plants	384	5,573	5,957
Multi-plant firms			
GHG Emissions	119.203	7.83	111.364***
	(10.951)	(0.261)	(1.257)
Employment	603.141	340.471	262.669***
	(25.367)	(2.386)	(3.563)
Carbon Intensity	0.381	0.037	0.345***
	(0.010)	(0.0008)	(0.004)
Coal Share	0.049	0.004	0.045***
	(0.002)	(0.0003)	(0.001))
Oil Share	0.149	0.134	0.015***
	(0.004)	(0.001)	(0.004)
Gas Share	0.432	0.361	0.071***
	(0.005)	(0.001)	(0.005)
Steam Share	0.014	0.010	0.003***
	(0.001)	(0.0004)	(0.001)
Observations	3,270	26,115	29,385
Plants	230	2,347	2,647
Single-plant firms			
GHG Emissions	81.276	4.168	75.63***
	(4.168)	(0.230)	(1.389)
Employment	479.652	242.962	236.690***
	(25.472)	(1.574)	(8.596)
Carbon Intensity	0.368	0.041	0.327***
	(0.018)	(0.001)	(0.007)
Coal Share	0.021	0.004	0.016***
	(0.002)	(0.0003)	(0.001)
Oil Share	0.168	0.147	0.020***
	(0.006)	(0.001)	(0.005)

Table 3: Descriptive statistics - ETS plants

	Mean	Std. Dev	Observations	No. of Plants	No. of Firms
Only ETS Multi-Plant Firms					
GHG Emissions	69.745	89.812	1,405	100	57
Employment	507.617	1225.922	1,405	100	57
Carbon Intensity	0.313	0.417	1,405	100	57
Coal Share	0.007	0.054	1,405	100	57
Oil Share	0.182	0.274	1,405	100	57
Gas Share	0.465	0.326	1,405	100	57
Steam Share	0.016	0.074	1,405	100	57
Multi-Plant ETS Firms with non-ETS plants					
GHG Emissions	156.462	823.738	1,865	130	76
Employment	675.103	1,596.274	1,865	130	76
Carbon Intensity	0.435	0.665	1,865	130	76
Coal Share	0.081	0.206	1,865	130	76
Oil Share	0.124	0.239	1,865	130	76
Gas Share	0.407	0.333	1,865	130	76
Steam Share	0.012	0.066	1,865	130	76
Single-Plant Firms					
GHG Emissions	81.276	191.365	2,107	154	154
Employment	479.652	1169.233	2,107	154	154
Carbon Intensity	0.368	0.857	2,107	154	154
Coal Share	0.021	0.101	2,107	154	154
Oil Share	0.168	0.168	2,107	154	154
Gas Share	0.484	0.326	2,107	154	154
Steam Share	0.018	0.083	2,107	154	154

Notes: Descriptive statistics are for all available years.

Figure 1: Proportion of firms that have multiple plants, by sector - All Plants

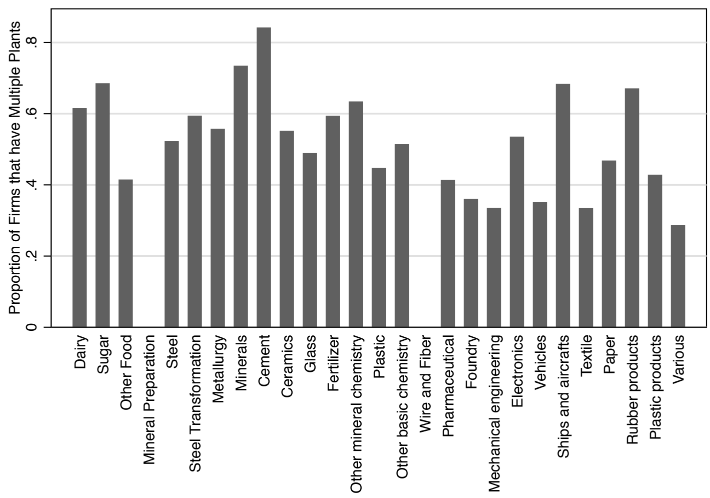

Figure 2: Proportion of ETS firms that have multiple plants, by sector - ETS Plants

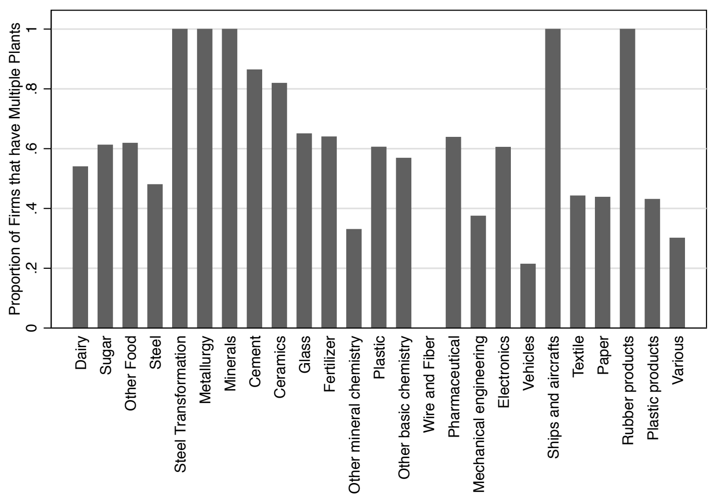

Figure 3: Proportion of non-ETS firms that have multiple plants, by sector - ETS Plants

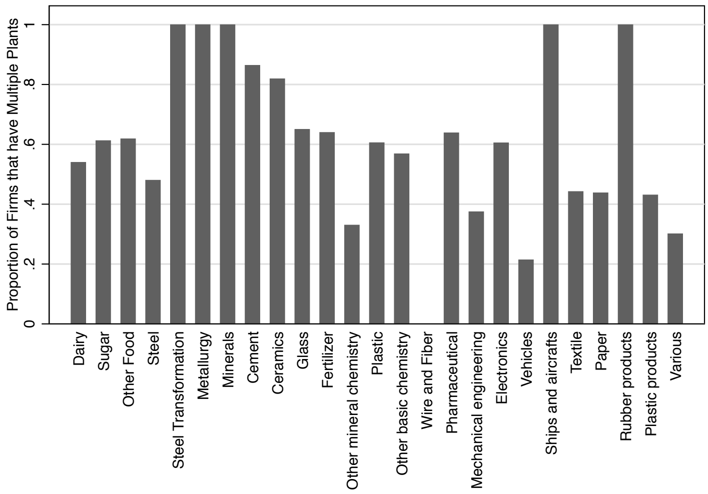

Plant Level Results

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

	(1) Base Year (2000)	(2) Base Year (2000)	(3) Base Year (2000)	(4) Base Year (2000)	(5) Base Year (2000)
Panel A: GHG emissions $\Delta ln(GHG\ Emissions)$					
Pre-Announcement SATT	-0.005	-0.006	-0.003	-0.003	-0.007
	(0.014)	(0.014)	(0.014)	(0.014)	(0.015)
Announcement Phase SATT	-0.002	-0.001	-0.002	-0.0009	0.002
	(0.012)	(0.012)	(0.012)	(0.013)	(0.014)
Phase I (2005-2007) SATT	-0.020	-0.030	-0.030	-0.032	-0.030
	(0.024)	(0.024)	(0.024)	(0.024)	(0.028)
Phase II (2008-2010) SATT	-0.135***	-0.151***	-0.155***	-0.168***	-0.198***
	(0.036)	(0.036)	(0.037)	(0.039)	(0.046)
Panel B: Employment $\Delta ln(Employment)$					
Pre-Announcement SATT	-0.008	-0.009	-0.006	-0.006	-0.011
	(0.009)	(0.009)	(0.009)	(0.009)	(0.010)
Announcement Phase SATT	-0.027***	-0.023**	-0.023**	-0.021**	-0.017
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Phase I (2005-2007) SATT	-0.010	-0.009	-0.012	-0.012	-0.015
	(0.023)	(0.023)	(0.023)	(0.023)	(0.026)
Phase II (2008-2010) SATT	-0.068**	-0.064**	-0.066**	-0.072**	-0.070**
	(0.030)	(0.030)	(0.029)	(0.030)	(0.033)
Panel C: Carbon Intensity $\Delta ln(GHG/Employment)$					
Pre-Announcement Phase SATT	0.003	0.003	0.003	0.003	0.003
	(0.014)	(0.014)	(0.014)	(0.014)	(0.015)
Announcement Phase SATT	0.024*	0.022	0.021	0.020	0.019
	(0.014)	(0.014)	(0.014)	(0.014)	(0.016)
Phase I (2005-2007) SATT	-0.010	-0.020	-0.018	-0.020	-0.014
	(0.026)	(0.025)	(0.026)	(0.027)	(0.030)
Phase II (2008-2010) SATT	-0.067*	-0.086**	-0.088***	-0.096***	-0.128***
	(0.034)	(0.034)	(0.034)	(0.035)	(0.041)
Observations	5,377	5,313	5,258	5,138	4,701
Clusters	287	285	283	283	270

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.

Base Year 2000

Figure 4: Difference in log(GHG) emissions and log(Employment) between treatment and comparison firms over time.

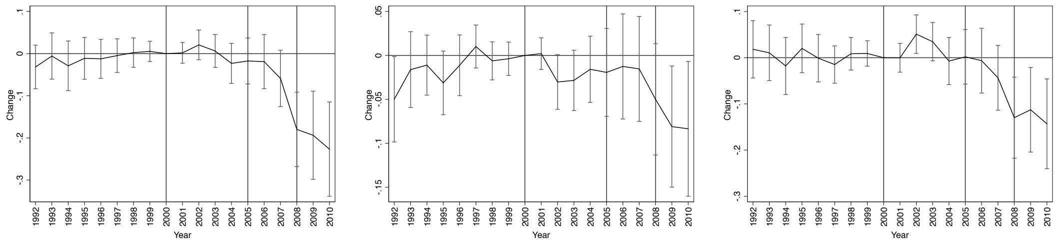

Base Year 2000 (maximum distance 100 percent)

Figure 5: Difference in log(GHG) emissions and log(Employment) between treatment and comparison firms over time.

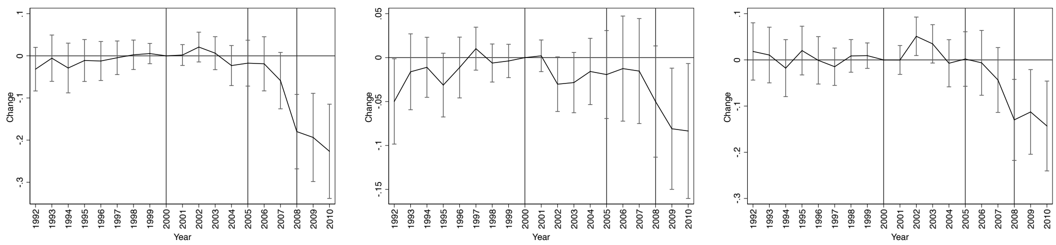

Base Year 2000 (maximum distance 50 percent)

Figure 6: Difference in log(GHG) emissions and log(Employment) between treatment and comparison firms over time.

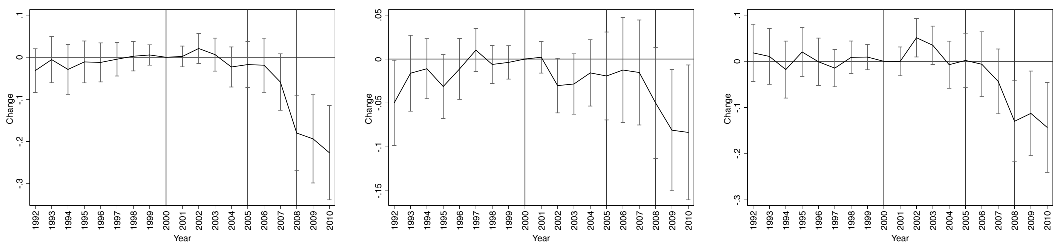

Base Year 2000 (maximum distance 10 percent)

Figure 7: Difference in log(GHG) emissions and log(Employment) between treatment and comparison firms over time.

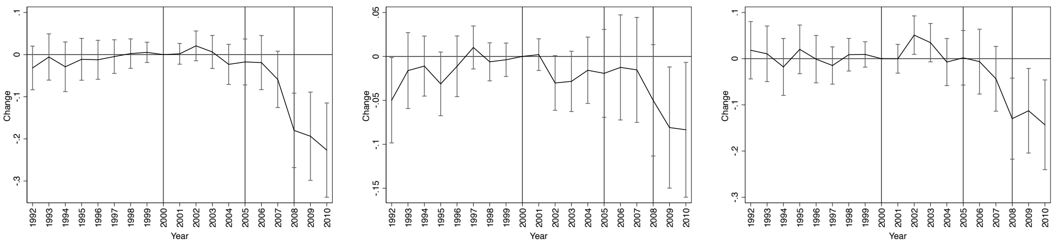

Firm Level Results

Table 5: Results from difference-in-difference propensity score matching estimators (firm level)

	(1) Base Year (2000)	(2) Base Year (2000)	(3) Base Year (2000)	(4) Base Year (2000)	(5) Base Year (2000)
Panel A: GHG emissions					
$\Delta ln(GHG\ Emissions)$					
Pre-Announcement SATT	0.005	0.004	0.004	0.004	0.004
	(0.015)	(0.015)	(0.015)	(0.015)	(0.016)
Announcement Phase SATT	0.014	0.015	0.014	0.018	0.021
	(0.015)	(0.015)	(0.015)	(0.015)	(0.016)
Phase I (2005-2007) SATT	0.012	0.012	0.012	0.001	0.030
	(0.030)	(0.031)	(0.031)	(0.031)	(0.034)
Phase II (2008-2010) SATT	-0.091**	-0.085**	-0.082**	-0.099**	-0.108**
	(0.038)	(0.039)	(0.039)	(0.039)	(0.047)
Panel B: Employment					
$\Delta ln(Employment)$					
Pre-Announcement SATT	-0.010	-0.011	-0.011	-0.011	-0.020
	(0.016)	(0.017)	(0.017)	(0.017)	(0.018)
Announcement Phase SATT	-0.040**	-0.039**	-0.038**	-0.037**	-0.036**
	(0.016)	(0.017)	(0.017)	(0.017)	(0.017)
Phase I (2005-2007) SATT	-0.026	-0.025	-0.027	-0.024	-0.012
	(0.027)	(0.027)	(0.027)	(0.028)	(0.029)
Phase II (2008-2010) SATT	-0.060	-0.050	-0.051	-0.056	-0.063
	(0.037)	(0.037)	(0.037)	(0.038)	(0.040)
Panel C: Carbon Intensity					
$\Delta ln(GHG/Employment)$					
Pre-Announcement Phase SATT	0.015	0.015	0.016	0.016	0.024
	(0.017)	(0.017)	(0.017)	(0.017)	(0.018)
Announcement Phase SATT	0.054***	0.054***	0.053***	0.056***	0.058***
	(0.018)	(0.018)	(0.018)	(0.019)	(0.019)
Phase I (2005-2007) SATT	0.039	0.038	0.039	0.026	0.042
	(0.029)	(0.029)	(0.030)	(0.030)	(0.033)
Phase II (2008-2010) SATT	-0.030	-0.034	-0.030	-0.042	-0.045
	(0.032)	(0.033)	(0.033)	(0.033)	(0.038)
Observations	3,893	3,822	3,765	3,665	3,239
Clusters	290	288	287	282	271

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.

Base Year 2000

Figure 8: Difference in log(GHG) emissions and log(Employment) between treatment and comparison firms over time.

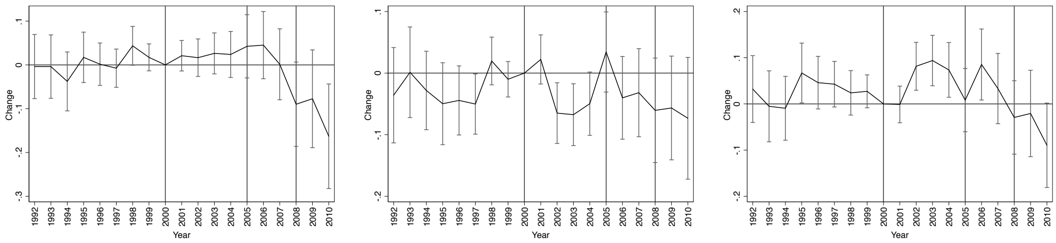

Base Year 2000 (maximum distance 100 percent)

Figure 9: Difference in log(GHG) emissions and log(Employment) between treatment and comparison firms over time.

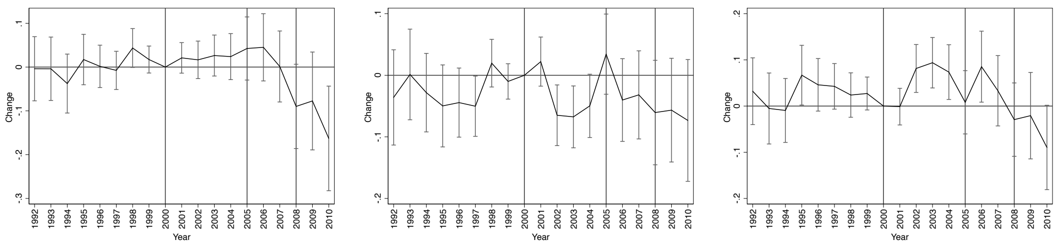

Base Year 2000 (maximum distance 50 percent)

Figure 10: Difference in log(GHG) emissions and log(Employment) between treatment and comparison firms over time.

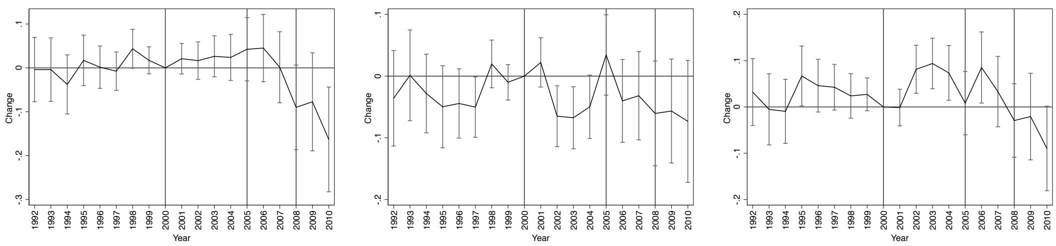

Base Year 2000 (maximum distance 10 percent)

Figure 11: Difference in log(GHG) emissions and log(Employment) between treatment and comparison firms over time.

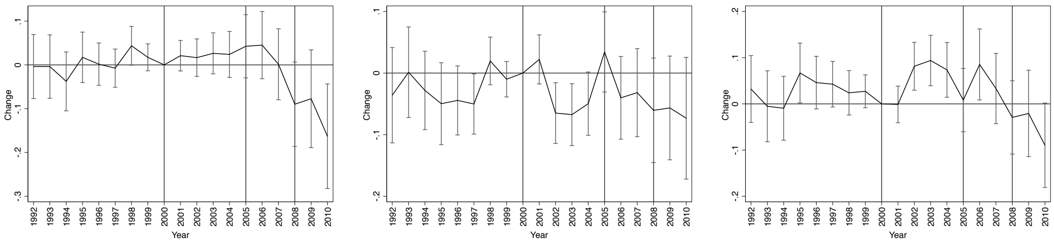

The Composition of the Firm

Table 6: Results from difference-in-difference propensity score matching estimators (plant level)

	(1) ETS (2000)	(2) ETS (2000)	(3) ETS (2000)	(4) ETS (2000)	(5) only ETS (2000)	(6) only ETS (2000)	(7) Non-ETS (2000)	(8) Non-ETS (2000)
Panel A: GHG emissions $\Delta ln(GHG\ Emissions)$								
Pre-Announcement SATT	0.012	0.020	0.019	0.031	0.009	0.009	0.045*	0.038
	(0.024)	(0.026)	(0.027)	(0.030)	(0.013)	(0.014)	(0.026)	(0.023)
Announcement Phase SATT	0.017	0.023	-0.005	0.009	0.010	0.017	-0.011	-0.012
	(0.028)	(0.034)	(0.023)	(0.018)	(0.014)	(0.015)	(0.029)	(0.061)
Phase I (2005-2007) SATT	0.008	0.027	-0.007	0.004	-0.011	-0.006	0.078	0.010
	(0.056)	(0.066)	(0.055)	(0.070)	(0.028)	(0.033)	(0.055)	(0.061)
Phase II (2008-2010) SATT	-0.111	-0.107	-0.104	-0.106	-0.122***	-0.166***	0.079	0.048
	(0.073)	(0.088)	(0.075)	(0.093)	(0.043)	(0.085)	(0.055)	(0.050)
Panel B: Employment $\Delta ln(Employment)$								
Pre-Announcement SATT	0.024	0.022	0.034*	0.033*	0.008	0.005	0.044	0.038
	(0.016)	(0.017)	(0.017)	(0.018)	(0.010)	(0.010)	(0.030)	(0.028)
Announcement Phase SATT	-0.050**	-0.041*	-0.071***	-0.046**	-0.028**	-0.022*	-0.027	-0.025
	(0.020)	(0.021)	(0.020)	(0.018)	(0.012)	(0.012)	(0.021)	(0.021)
Phase I (2005-2007) SATT	-0.033	-0.018	-0.037	0.016	-0.013	-0.014	0.042	0.041
	(0.041)	(0.044)	(0.041)	(0.045)	(0.024)	(0.028)	(0.047)	(0.049)
Phase II (2008-2010) SATT	-0.100*	-0.057	-0.095	-0.059	-0.067**	-0.066*	0.038	0.037
	(0.058)	(0.063)	(0.060)	(0.065)	(0.044)	(0.040)	(0.037)	(0.039)
Panel C: Carbon Intensity $\Delta ln(GHG/Employment)$								
Pre-Announcement SATT	-0.012	-0.002	-0.014	-0.002	0.001	0.003	0.001	-0.00007
	(0.026)	(0.028)	(0.027)	(0.029)	(0.014)	(0.015)	(0.024)	(0.026)
Announcement Phase SATT	0.067**	0.065*	0.066**	0.055*	0.038**	0.039**	0.015	0.012
	(0.028)	(0.034)	(0.026)	(0.033)	(0.015)	(0.016)	(0.028)	(0.027)
Phase I (2005-2007) SATT	0.041	0.045	0.029	0.020	0.001	0.008	0.035	-0.030
	(0.047)	(0.056)	(0.048)	(0.059)	(0.027)	(0.032)	(0.044)	(0.043)
Phase II (2008-2010) SATT	-0.011	-0.050	-0.009	-0.047	-0.054	-0.099**	0.040	-0.010
	(0.061)	(0.071)	(0.062)	(0.076)	(0.036)	(0.044)	(0.051)	(0.052)
Observations	1,355	1,166	1,350	1,148	4,163	3,665	1,625	1,491
Clusters	65	61	65	61	245	233	64	63
Non-ETS plants in ETS firms included in sample	Yes	Yes	No	No	N/A	N/A	Treatment	Treatment

Notes: Treatment plants are matched to control plants based on carbon intensity, with exact matching on sectors at the 2 digit level. Columns 1, 3, 5, and 7 restricts the distance to 50%. Columns 2, 4, 6, and 8 restricts the distance to 10%. Significance levels are indicated as * 0.10 ** 0.05 *** 0.001. Robust Standard errors, clustered at the firm level, are in parentheses.

The Composition of Emissions

Table 7: Results from difference-in-difference simple average matching estimators (plant level)

	(1) (2000) Coal	(2) (2000) Coal	(3) (2000) Oil	(4) (2000) Oil	(5) (2000) Gas	(6) (2000) Gas	(7) (2000) Steam	(8) (2000) Steam
Pre-Announcement SATT	0.004	0.004	0.0001	0.0001	0.002	0.001	-0.003	-0.003
	(0.003)	(0.003)	(0.007)	(0.007)	(0.007)	(0.008)	(0.002)	(0.003)
Announcement Phase SATT	-0.006	-0.010**	-0.014***	-0.012**	0.020***	0.023***	0.002	0.002
	(0.004)	(0.004)	(0.005)	(0.005)	(0.006)	(0.006)	(0.003)	(0.003)
Phase I (2005-2007) SATT	-0.003	-0.008	-0.010	-0.007	0.009	0.013	0.003	0.006
	(0.006)	(0.007)	(0.010)	(0.012)	(0.012)	(0.015)	(0.005)	(0.005)
Phase II (2008-2010) SATT	-0.018**	-0.023**	-0.016	-0.016	0.020	0.021	0.005	0.006
	(0.008)	(0.009)	(0.012)	(0.015)	(0.015)	(0.019)	(0.005)	(0.005)
Observations	5,258	4,701	5,258	4,701	5,258	4,701	5,258	4,701
Clusters	283	270	283	270	283	270	283	270

Notes: Treatment plants are matched to control plants based on carbon intensity, with exact matching on sectors at the 2 digit level. Columns 1, 3, and 5 restricts the distance to 50%. Columns 2, 4, and 6 restricts the distance to 10%. Significance levels are indicated as * 0.10 ** 0.05 *** 0.001. Robust Standard errors, clustered at the firm level, are in parentheses.

Appendix 1 - 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 12: Distribution of pre-treatment log(Carbon Intensity) for ETS plants and non-ETS plants: before matching and after matching

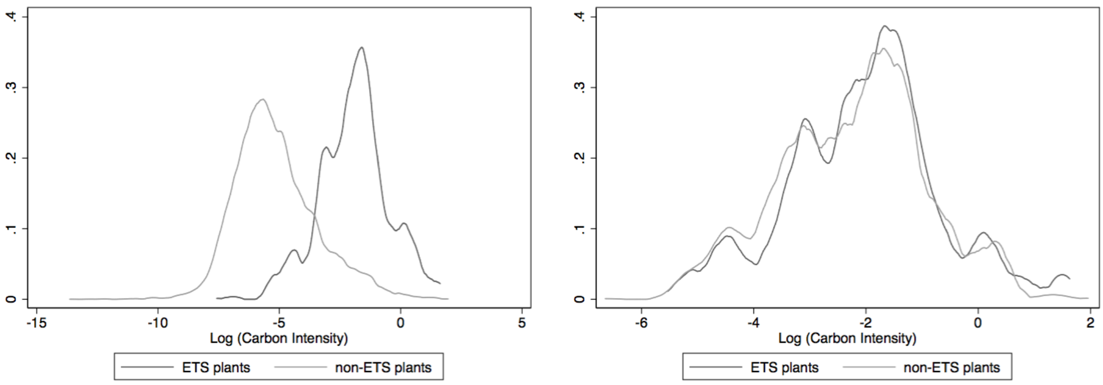

Figure 13: QQ plot of pre-treatment log(Carbon Intensity) for ETS plants against pre-treatment log(Carbon Intensity) for non-ETS firms before matching and after matching

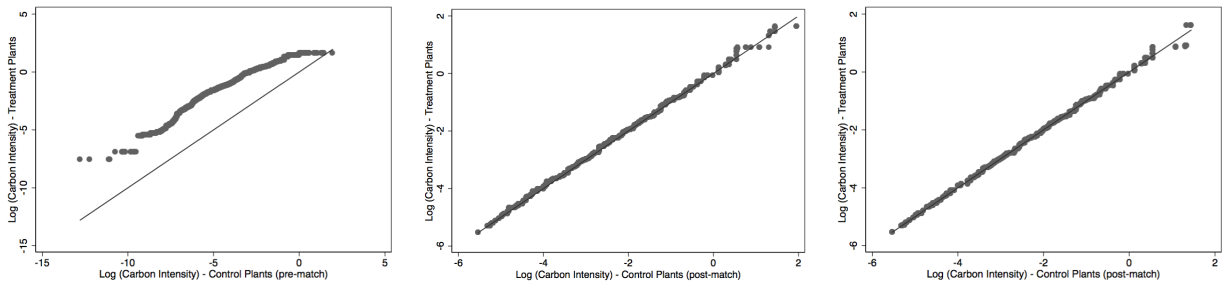

By comparing the plots in figure 12 we can see the improvements that matching has made in the distributions of the treatment and control groups. The QQ plots in figure 13 further demonstrate the improvements that matching can make. Panel A in figure 13 shows the probability distributions of the treatment 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).

As an additional check, we examine the distribution of other pre-treatment characteristics before and after matching that were not included in the covariates used for matching: log(GHG) and log(Employment).

Figure 14: Distribution of pre-treatment log(GHG Emissions) for ETS firms and non-ETS firms: before matching and after matching.

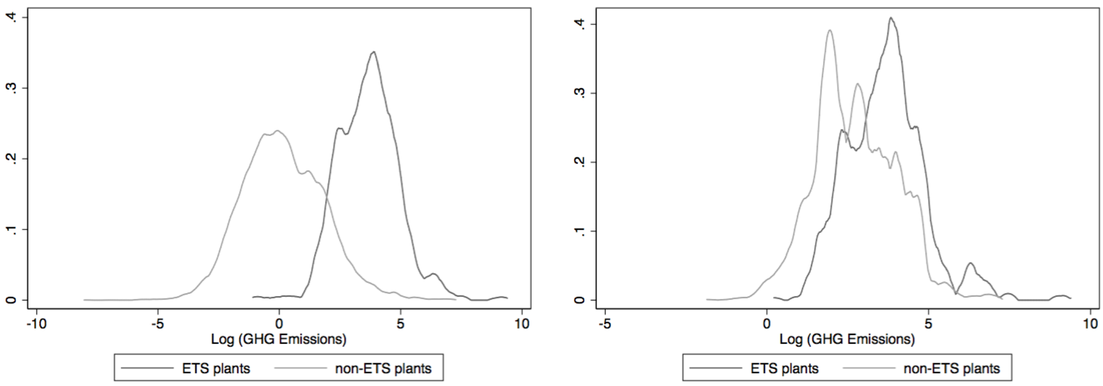

Figure 15: QQ plot of pre-treatment log(GHG Emissions) for ETS firms against pre-treatment log(GHG Emissions) for non-ETS firms before matching and after matching

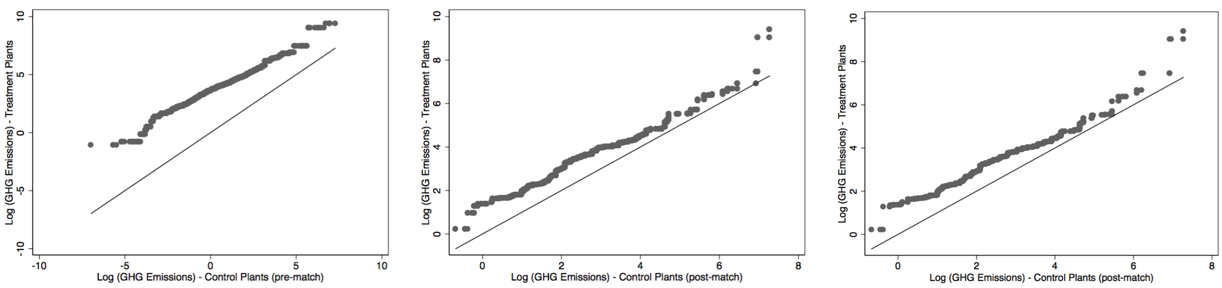

Figures 14 and 15 indicate that there has been a significant improvement in the balance between the treatment and comparison groups for greenhouse gas emissions even without matching on this covariate. Understandably, this balance isn't as precise compared with carbon intensity, however, it is a vast improvement on the unconditional difference.

Figure 16: Distribution of pre-treatment log(Employment) for ETS firms and non-ETS firms: before matching and after matching

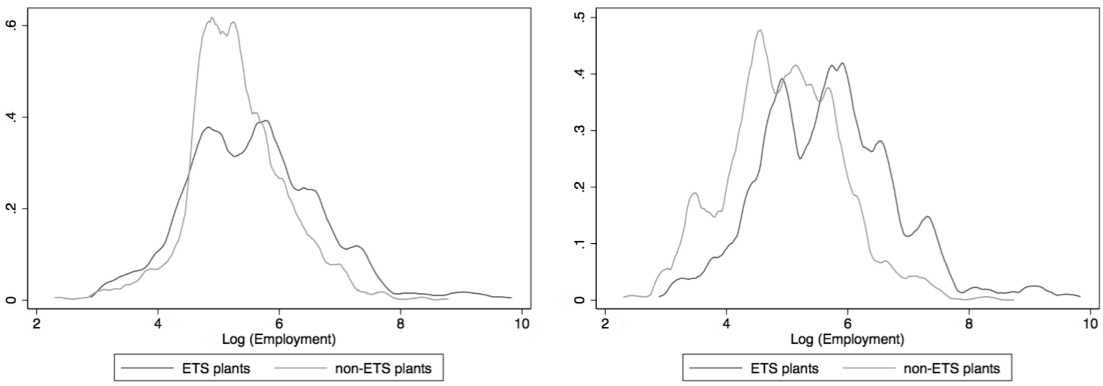

Figure 17: QQ plot of pre-treatment log(Employment) for ETS firms against pre-treatment log(Employment) for non-ETS firms before matching and after matching

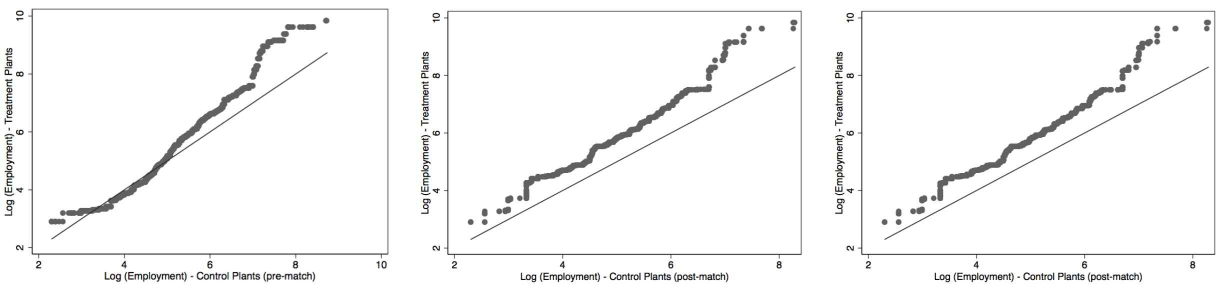

Figures 16 and 17 examine the balance between pre-treatment employment. Again, we see that the distribution is much improved compared to the unconditional difference. The QQ plots in figure 17 demonstrate this as well.

The QQ plots in panel B of figure 11 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%, 25% and 10%.

Figure 18: Nearest neighbours included when the maximum distance between treatment and control is restricted to 100%

Figure 19: Nearest neighbours included when the maximum distance between treatment and control is restricted to 50%

Figure 20: Nearest neighbours included when the maximum distance between treatment and control is restricted to 25%

Figure 21: Nearest neighbours included when the maximum distance between treatment and control is restricted to 10%