

The EU ETS's time-varying impact on Competitiveness and Investments: Evidence from Dutch Manufacturing

Leon Bremer* Konstantin Sommer†

January 31, 2022

Abstract: We study the effects of the EU ETS on competitiveness and investment decisions of Dutch manufacturing firms. We pay close attention to treatment effect heterogeneity by studying the heterogeneous effects for firms starting in different ETS phases. We use microdata from Statistics Netherlands (CBS) to apply two difference-in-differences (DiD) estimators. We employ both a two-way fixed effects regression, as well as a newer and more flexible DiD method that allows for multiple treatment periods, aligning with the phases of the EU ETS. We find no reduction in competitiveness in regulated firms, but in fact observe that companies that start in phase 1 both become larger and invest more. While the positive effect on competitiveness peaks in the mid 2010s, the effect on investments increases over time. For firms starting in later phases, we do not find any statistically significant effect of ETS participation.

*PhD Candidate at Vrije Universiteit Amsterdam – L.Bremer@VU.NL.

†PhD Candidate at University of Amsterdam – K.H.L.Sommer2@UvA.NL. The authors are grateful for funding by the “A Sustainable Future (ASF)” platform from the University of Amsterdam awarded to Konstantin Sommer.

1 Introduction

This paper studies the effects of the European Union Emissions Trading System (EU ETS) on both the competitiveness as well as the adoption of new technologies of regulated manufacturing firms. In our analysis we carefully disentangle the effects of firms that become regulated at different moments in time as well as the dynamics of the effects over time.

The EU ETS is the world's largest cap-and-trade system, aiming to reduce the EU's emissions in the manufacturing, energy and heat sector, as well as in intra-EU aviation. It was implemented in 2005 and caps the annual amount of emissions within the EU. Large emitters have to commit one allowance for each emitted ton of carbon at the end of each year. These allowances can be traded on financial markets, thereby establishing a price for carbon. ETS regulation has changed over different phases, in which both new installations became regulated and in which regulation and its stringency were amended. As we show, firms regulated in phase 1 are the largest and most energy-intensive firms in the Netherlands, while firms regulated in later phases are considerably less energy-intensive and mostly smaller, though still larger than the average Dutch firm.

Since the beginning of the ETS, policy makers and industry representatives have been concerned about the problem of (carbon) leakage. This would occur if EU producers face such high carbon prices that they would lose competitiveness to their non-EU competitors. This could lead to a loss of jobs in Europe accompanied by an increase in emissions elsewhere. This would be undesirable from an EU welfare as well as from an environmental point of view. Meanwhile, stricter regulation could also incentivize firms to invest in new technologies and gain competitiveness in the long run, a hypothesis in line with Porter and Van der Linde (1995).

While some studies estimate the effects of the ETS on both competitiveness and investment behavior, none of these analyze the underlying treatment effect heterogeneity of the policy, a topic that has recently received great attention in econometric literature. Furthermore, few studies cover the more recent years in which allowance prices are on the rise. This paper tries to close these gaps. We show that firms entering regulation at different points in time are substantially different from each other, making pooled difference-in-differences (DiD) estimates potentially misleading. We estimate these heterogeneities and compare the results to a more standard approach using data up to 2019. Further, this paper is the first to conduct such a study using Dutch microdata.

We are able to use detailed firm-level microdata from Statistics Netherlands (in Dutch: CBS), the Dutch national statistics agency, and link those to the European Union Transaction Log (EUTL) for information on regulated ETS firms. Besides the more classic DiD implementation through matched two-way fixed effects

(TWFE), we employ a recent, more flexible DiD method developed in Callaway and Sant’Anna (2021) that allows for multiple treatment periods and heterogeneity over event time. This more flexible approach allows us to disentangle the average treatment effect on the treated (ATT) over different groups, time periods and event time. This offers ways to break down the effects of the EU ETS between its phases and between its participants.

We find that regulated firms tend to increase the number of employees. This finding is driven by the group first regulated in phase 1. Firms that start treatment in later periods do not experience significant losses or gains in competitiveness. The same conclusion can be drawn for value added, where the group regulated in phase 1 of the EU ETS increased its average value added, but other groups did not significantly respond to the regulation. For both dependent variables this effect is largest in the mid-2010s and seems to be reverting to zero in the last years of our sample. This indicates that the large phase 1 group, first regulated in 2005, if anything increased their competitiveness, while other firms have seen little to no impact on their competitiveness. These results stem from the analysis with the Callaway and Sant’Anna (2021) approach. Using the more regular TWFE estimator leads to insignificant coefficient estimates, which might be driven by the pooling of underlying heterogeneities.

We find evidence that the regulation under the EU ETS caused investments in fixed assets to increase. The matched TWFE method does not result in statistically significant coefficient estimates for any of the first three EU ETS phases, but when disentangling these effects over the group, time and event time dimensions, we find that the group of firms regulated first in phase 1 responded to the regulation by a statistically significant increase in fixed assets investments. There further seems to be an increase in the effect over time for this group. These findings highlight the importance of incorporating time and group heterogeneity into DiD analyses.

The paper continues as follows. [Section 2](#) discusses the related literature. The data and policy background are discussed in [section 3](#). The methodology and results are presented in [section 4](#) and [section 5](#) respectively. [Section 6](#) concludes.

2 Related literature

There are several studies on the effects of the ETS on the competitiveness of regulated firms as well as potentially resulting carbon leakage and related technology adoption. Some studies use administrative firm-level data in other countries than the Netherlands and apply comparable difference-in-differences (DiD) studies. Other studies utilize a larger set of EU ETS firms combined with publicly available data sets (e.g. Calel & Dechezleprêtre, 2016). Underlying all studies is the complexity of finding appropriate control firms that are unregulated but

sufficiently similar, e.g. through matching, in order to draw causal conclusions.

Most studies relying on matching estimate the treatment effect on the treated (ATT) by using the semi-parametric estimator of Heckman et al. (1997). In this framework Wagner and Petrick (2014) and Jaraite-Kažukauske and Di Maria (2016) cannot find negative effects of the ETS on productivity and employment for Germany and Lithuania, respectively. Löschel et al. (2019) additionally use a two-way fixed effects (TWFE) setting to analyze the ETS’s effect on productivity in Germany. The authors interestingly find significant positive effects on productivity using the Heckman-style estimator, but not in the regression estimation. Marin et al. (2018), using non-administrative micro data from Bureau van Dijk for a larger set of countries, also do not find negative effects on economic performance, but do find an increase in labor productivity. The only study that establishes some form of competitiveness loss is by Wagner et al. (2014), who find a significant reduction in employment, starting from phase 2 onwards.

All of these studies, however, only use data for the first phase and some years into the second phase. Since the stringency of the ETS increased significantly in the second and third phases, these studies potentially miss the largest effects that the ETS has had. Dechezleprêtre et al. (2019) use data on multinational firms up to 2014 and analyze carbon shifting within these firms, again without finding much evidence of leakage. A comparable study to ours that also looks at phase 3 is by Klemetsen et al. (2020) and analyzes firms in Norway. The authors use a regression setting, differentiating the effects between phases, but not between companies starting in different phases, and find a slight increase in productivity in phase 2, but no significant effect in the other phases.

In a literature survey, Verde (2020) comes to the conclusion that there is no convincing evidence of leakage and losses in competitiveness due to the ETS yet. The authors also highlight that this might be due to the short time span covered in almost all studies and point to the importance of analyzing more long-term indicators like investments. This study is trying to address both of these gaps.

Further, to address such time-varying heterogeneity in the treatment effects, the econometric design needs to be appropriate. There is recent discussion of DiD and its approximation using TWFE (see e.g. de Chaisemartin & D’Haultfoeuille, 2022). Callaway and Sant’Anna (2021) offer a more flexible alternative to TWFE. In this paper we will use their estimator as well.

Our contribution to the previous work is threefold. First, we add analytically to the debate about causal effects of the ETS, by showing that previous studies might have tried to estimate homogeneous treatment effects in cases where homogeneity seems unlikely. For this we extend the analysis to both a more flexible fixed effect estimator and to a more detailed semi-parametric estimator, developed in a recent stream of econometric literature. Second, we benefit from longer

time series, allowing us to estimate the later phase’s effects. Third, we are able to use detailed administrative data on Dutch firms. We have access to data on investments, on top of the more classic indicators for firm performance and competitiveness. Additionally, the Netherlands are due to their export orientation and rather energy-intensive industrial structure a country in which competitiveness effects might be felt strongly.

3 Data and policy background

3.1 EU ETS policy background

The EU ETS regulates installations, which we will also refer to as plants. Each of these plants is registered under one account holder at a time in the European Union Transaction Log (EUTL). The amount of regulated installations and account holders in the Netherlands can be found in Figure 1. After its initial implementation in 2005 the ETS has been largely revised 3 times when new phases came into effect, in 2008, 2013 and 2021. Most of these revisions aimed at making the system more restrictive and effective.

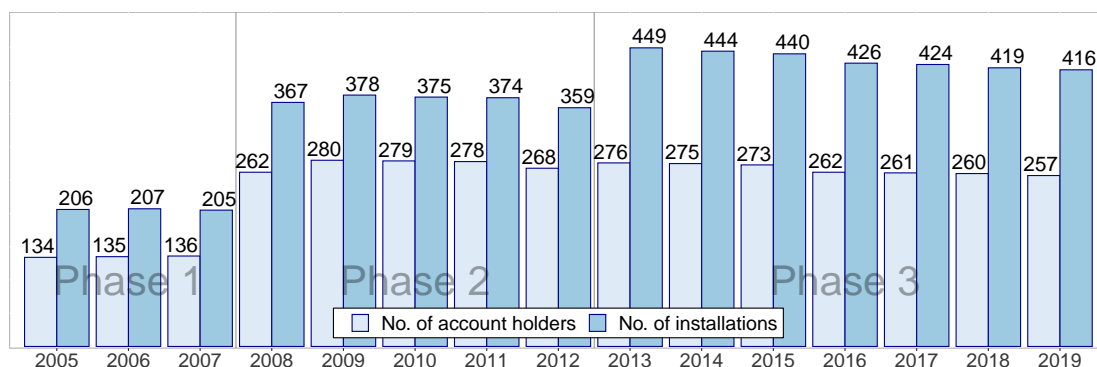


Figure 1: Number of active account holders and installations under the EU ETS over time in the Netherlands.

In phase 1 (2005-2007) allowances were handed out so plentiful that their price dropped to zero towards the end of the phase. In the Netherlands actual emissions were almost 15 percent below the number of allocated allowances (Ellerman & Buchner, 2008). Phase 2 (2008-2012) added nitrous oxide as a greenhouse gas and increased the penalty for non-compliance from €40 to €100 per tonne of CO₂ equivalent. The amount of regulated installations within the Netherlands increased from 205 to 367 (see Figure 1), mostly because in Phase 1 150 Dutch installations

were excluded from the ETS.¹

Even more greenhouse gases were added in Phase 3 (2013-2020). Also, the default allowance allocation method switched from grandfathering to auctioning. During phase 3 the Market Stability Reserve (MSR) entered into force, backloading new allowances and cancelling excess allowances if needed, adjusting for sustained periods of low demand. Further, manufacturing sectors in the aluminium and chemicals production were added to the coverage. This did not change the number of regulated account holders much, but it significantly increased the number of regulated plants (see [Figure 1](#)). Arguably more plants of the same owners were regulated in phase 3. Phase 4 (2021-2030) mainly sped up the rate at which the cap decreases over time and it strengthened the MSR.²

The changing degree of stringency is also reflected in the allowance price path, as depicted in [Figure 2](#). Prices decreased to zero at the end of phase 1, then started around €20 in phase 2, but stayed around only €10 for several years. Even though economists argue about the optimal price of carbon, such low prices have almost uniformly been deemed as too low to have the intended impact. Prices have started to increase since 2017/2018 and have nearly reached €90 at the start of 2022, making the ETS far more restrictive in recent years.

To identify the effects of the ETS, we use the fact that not all manufacturing firms in the EU are regulated under the ETS. Regulation is on the plant level and there are mainly two criteria for plant inclusion in the ETS, either (1) through exceeding a certain sector-specific threshold related to energy input or production capacity or (2) through incorporating specific processes that imply automatic regulation.³ This implies that one can attempt to find comparable control firms for each treated firm, that are both active in comparable production processes and are comparable in terms of economic variables like size, employment characteristics and or energy input. We also use the differentiated treatment timing to estimate the effects on different cohorts of firms that were regulated.

3.2 EUTL and Dutch microdata

The data for this project comes from two main sources. First, the European Union Transactions Log (EUTL) data is accessed through [EUETS.INFO](#), a free service that provides cleaned data from the EUTL. Second, Dutch firm-level data

¹The following decisions by the European Commission (EC) provide further details of the phase 1 exemptions for Dutch installations. In October 2004 the EC exempted 93 installations and in March 2005 the EC exempted a further 57 installations (European Commission, [2004](#), [2005](#)).

²Please refer to the [European Commission's webpage](#) for more details.

³For a detailed overview see Annex I of European Parliament, Council of the European Union ([2003](#)).




Figure 2: *The EU ETS’s allowance price in Euros per tonne of CO₂ equivalent. These are day closing prices for its futures contracts. The data are accessed through FactSet. The futures montage ECF00-NDEX is plotted in solid blue. The December 2007 futures price for phase 1 allowances is plotted as a dotted orange line. These allowances were not transferable to later phases. The phase 1 data come from the European Environment Agency.*

is accessed through the microdata services of Statistics Netherlands (in Dutch: CBS). This data contain firm-level information on economic activity of almost the entire population of Dutch firms with more than 50 employees.

The data collected from the EUTL contain information on the free allocations of allowances (EUAs), verified emissions, allowances surrendered, and the use of international credits, both by installation and account holder. The data are organized in an unbalanced panel spanning the years 2005-2019 and a total of 439 unique account holders, owning 598 installations.

Account holders can potentially own several regulated plants and are registered under a national identification number.⁴ We link these data to the administrative firm-level data of CBS, the Dutch statistical agency. CBS data are not publicly accessible and are anonymized, which is why after linking the EUTL data to CBS’s data we are not able to identify individual firms anymore.

⁴As installations are assets, they can be purchased from or transferred to other firms. Such changes of ownership are not perfectly captured by the data. Many installations do not change ownership between EU ETS phases in our data, but for the ones where it does change, we manually looked up the date of ownership change using online public sources. Sources can be online news articles or websites that provide information about ownership structures. The list of manually assigned ownership changes and their respective source is available upon request.

The CBS data contain a host of information on the firm level. It contains information like number of employees, balance sheet information, as well as investment data and international trade flows. This study is restricted in scope to manufacturing firms and relies on more than 40 thousand firms over a time span of 20 years. To deflate monetary variables, we use Eurostat’s industry producer price index for the Netherlands.

CBS data contain firm-level data, and not plant-level data. Within CBS, several chamber of commerce numbers can comprise a business unit, a construct defined by CBS and further explained in [Appendix A.3](#). We will from here on refer to these business units as “firms”. We are able to link the account holders from the EUTL with the business units in the CBS data.

As a business unit can comprise of multiple account holders and plants, it can be the case that a business unit is regulated only through one plant or through multiple plants. We do not make a distinction here and consider each business unit (firm) as regulated if it owns at least one regulated plant. Our level of analysis is on this business unit level, referred to as the firm level.

3.3 Descriptive statistics

[Figure 3](#) shows the development of the average treated and non-treated firm over our sample time for energy expenditure, employment, value added and investments in fixed assets. The plot shows averages for ETS firms over the three phases compared to both the entire sample (unmatched) and a more similar sample (matched). The matched sample will function as control group in the later analysis. The matching procedure is described in [section 4](#).

Most strikingly, one can see that firms regulated in 2005 are by far the largest energy consumers among all Dutch firms. Which makes sense, since they were also chosen to be regulated first. In terms of employment and value added, treated firms seem to be more similar between phases, even though the latest treated firms appear to be the smallest ones.

In comparison to the rest of the Dutch firms, ETS firms are larger and more energy intensive, which is again unsurprising and which implies that we have to take care of these fundamental differences. For both our estimation approaches, however, we do not require firms in control and treatment groups to be the same in terms of their outcome levels, but in terms of their trends. From these plots one can only draw anecdotal support, but it does not seem as if treated firms behave significantly different from control firms in the pre-treatment period. For the treated firms one can however also not see significant changes in the trends after their treatment starts, which suggests no large treatment effects.

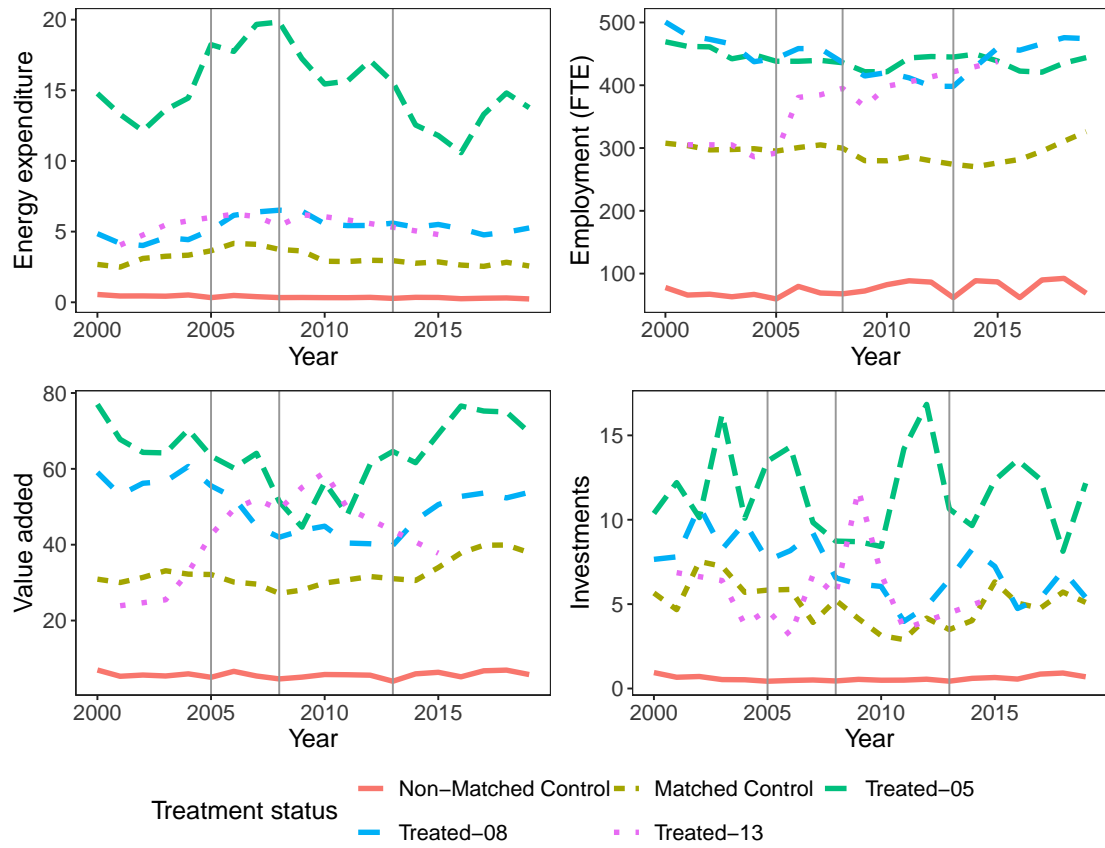


Figure 3: *Descriptive statistics for energy expenses, employment, value added and investments in fixed assets for regulated and non-regulated firms. All monetary variables are expressed in 2015 Million Euros.*

4 Methodology

To evaluate the causal effects of the EU ETS climate policies, we apply two empirical policy evaluation methods. The key in these methods is to use detailed microdata on observed firms to compare the outcomes of treated firms, i.e. firms receiving regulation, to the outcomes of *comparable* control firms.

In general two main steps can be identified in this evaluation process, namely (1) matching or weighting, in which we score firms across treatment status based on their similarity, and (2) comparison, in which we either regress our outcome variable on treatment status or take difference in outcome variables, where observations are weighted by the outcome of step 1.

Both empirical methods rely on weighting and a DiD design. The first method is a matched two-way fixed effects (TWFE) regression and the second method is a less structured DiD design suggested by Callaway and Sant’Anna (2021). The latter method has several advantages. Amongst other things it allows us to disentangle the results per treatment group, and over time and event time.

These methods should give us insight in the ETS’s effect on (1) competitiveness and (2) investments. Both the number of employees and value added are used here as proxies for competitiveness. If the ETS negatively affects competitiveness of regulated firms, we expect to see a decrease in the number of employees and firms’ value added. If firms respond to the EU ETS by updating their assets, we expect to see this an increase in the total investments in fixed assets.

4.1 Matched TWFE method

As mentioned above, our first method relies on matching and regression. We break these two steps up in the following two sections. The first section explains the matching that provides the weights, and the second section presents the details and form of the regression.

4.1.1 Matching

The goal of matching is to select similar observations across treatment status from the data. In general a matching algorithm provides a similarity score between each pair of observations in the sample data. If provided with n observations, the matching outcome matrix M has dimensions $n \times n$. For our application the pair information is dropped and only those observations with a high enough similarity scores are kept. Observations in the non-treated group that do not have a high enough similarity score with a treated observation are thus dropped from our sample. This way matching boils down to sample selection.

Our matching algorithm is presented in [Algorithm 1](#). This algorithm does not return the underlying matching matrix, but it returns connections between matched treated units, matching year, and matched control units.

Algorithm 1: *Matching*

1. Select treatment period
 - (a) Take treatment period $T \in T^p$, where T^p is the set of treatment periods, i.e. the years 2005, 2008 and 2013 for phase 1, phase 2 and phase 3 (p) in the EU ETS respectively.
2. Select observations to be potentially matched
 - (a) From the ever-treated EU ETS firms, select only those observations that are first regulated in phase p . Keep all observations from the never-treated group.
 - (b) Only keep units that are observed for all of the years in $(T - pre, T + post)$, where we set $pre = 2$ and $post = 4$. This guarantees that resulting matches can be observed around the treatment period.
 - (c) Select only the observations at $T - pre$, dropping the panel structure. This year will be the pre-treatment matching period.
3. Similarity scoring and match decision
 - (a) Measure the Mahalanobis distance between all observations in the selected sample across treatment status for the variables X^m .⁵ X^m are the matching variables for which we take the number of employees, turnover, wage expenses, energy expenses, and value added. We also restrict matches to be only within a 2-digit sector code. Matches across sectors are not allowed.
 - (b) For each treated unit collect the H closest neighbors based on the Mahalanobis distance. We opt for $H = 5$ and we do allow for replacement. We also allow for ties, meaning ties are not randomly broken but rather all are included in the result. For the implementation of this step and the previous step we leverage on the `Matching` package's `Match` function in R.

⁵The Mahalanobis distance between treated (T) unit A's covariate vector x_A and control (C) unit B's covariate vector x_B is given by $d(A, B) = \sqrt{(x_A^T - \mu^T)S^{-1}(x_B^C - \mu^C)}$, where S is the variance-covariance matrix between x^T and x^C and where the μ s are the means of their respective series. Note that this distance measure is like a variance-corrected normalized Euclidean distance.

4. Match correction

- (a) To avoid matches that are too far from each other, we run a post-match filter. Since the number of nearest neighbors is fixed, the distance between matches across treated units can otherwise vary substantially. The post-match filter looks up the values of X^m for the treated and its matched units and filters out any matched values that are p percent larger or smaller than the treated unit's values. We set $p = 80\%$.

5. Store matching outcome

- (a) Remaining matches are stored under matching year $T - pre$.

6. Next treatment period

- (a) If not all treatment periods in T^p are covered yet, select the next value in T^p and repeat the algorithm from step 2.

These matching outcomes are used to select the sample for our TWFE regression. All observations are kept that have an identifier in the matching outcome, either in the treatment or the control group. This effectively is a special form of weighting, as the weights are either 1 (for the matched) or 0 (for the non-matched).

4.1.2 TWFE regression

Taking the above matching outcome as given, we can perform a DiD regression closely related to that of Klemetsen et al. (2020). This gives us the following two-way fixed effects (TWFE) regression

$$y_{i,t} = \sum_{p \in \{1,2,3\}} ETS^p \delta^p + T^p \rho^p + ETS^p \times T^p \alpha^p + X_{i,t} \beta + \gamma_i + \varepsilon_{i,t} \quad (4.1)$$

where ETS^p is an indicator of treatment under phase p of the EU ETS, T^p an indicator of time in EU ETS phase p , X is a vector of covariates, γ are fixed effects, ε is the error term, i and t indicate the firm and year respectively, p runs over the three phases of the EU ETS, and y is the dependent variable. Note that we assign each ETS firm to the phase in which they are first regulated. Similarly $T^p = 1$ if and only if year t falls in the years in which phase p is active.

The dependent variable y is either the number of employees or value added when interested in the competitiveness effects from the EU ETS, and investments in fixed assets when interested in the investments response. The control variables in X are turnover, value added, number of employees, total wages, investments

in fixed assets, and energy expenditures. When relevant, variables in y and X are deflated using EUROSTAT’s inflation figures for the Netherlands. And all variables in y and X enter the regression in log form.

4.2 CS2021’s DiD estimator

4.2.1 Potential problems with the TWFE approach

In addition to the classic TWFE estimator, we follow a recent stream of econometric literature dealing with difference-in-differences (DiD) estimators in a setting with multiple treatment periods. This literature focuses on the potential bias in TWFE estimators applied to such settings (see e.g. Daw & Hatfield, 2018; de Chaisemartin & D’Haultfoeuille, 2022; Goodman-Bacon, 2021). The key problem of TWFE is that the derived estimator for the ATT is a weighted average over the ATTs of the different treatment groups at different times. The estimator does thus only give you a clearly interpretable ATT if treatment effects are constant both over time and between treatment groups.

Our TWFE setting for example implicitly assumes that the effect of an ETS phase is homogeneous for all firms that are treated within that phase. For this, one has to keep in mind that some of the firms that are treated in phase 2 and 3 have already been regulated for one or two phases, while others just start regulation in that phase.

There are thus two main reasons for why the assumptions underlying TWFE estimation might not hold. The first is that the treatment effect could be dynamic, e.g. when earlier treated firms respond differently from newly treated firms. The second is that newly regulated firms might respond differently to regulation than firms for which the new phase only presents an adaptation of the regulation. The treatment effect for phase 2 would then for example be the weighted average of firms that just started treatment, and would thus be new to regulation, and firms from phase 1 whose effect is driven by a longer adjustment to the regulation, because they have been exposed to some kind of treatment before, and who might not face a shock from being newly regulated.

An additional problem with the TWFE estimators lies in the inclusion of time varying control variables. Especially in the ETS it is likely that the effect of the ETS affects both the variable of interest as well as the included controls. It is hard to find a control variable that is both time varying but unaffected by the ETS treatment.

4.2.2 The estimator

To address all of these issues, we make use of the estimator developed by Callaway and Sant’Anna (2021). Its advantage lies in the fact that it estimates ATTs for each treatment group – the group of firms starting treatment in the same phase – and at each year after treatment. But it also provides intuitive aggregations of those estimates. These aggregations allow us to restrict the heterogeneity, by making explicit what the underlying assumptions for these restrictions are.

The estimator is in essence an application of the doubly robust DiD estimator developed in Sant’Anna and Zhao (2020). It pays close attention to the conditioning on covariates, combining both inverse probability weighting (Abadie, 2005) as well as outcome regression adjustment (Heckman et al., 1997). The latter is also frequently used in adjusted versions in comparable ETS papers like Martin et al. (2014) or Löschel et al. (2019).

The estimator for each treatment group $g \in \{2005, 2008, 2013\}$ and year t is a common (semi-parametric) average treatment effect DiD estimator. In such, we are comparing the outcome of each treatment group (cohorts starting in different ETS phases) in year t to its outcome in the (cohort-specific) base year, b , and to that of the weighted average of the respective control group for this cohort. For this weighting, both inverse probability weighting and outcome regression techniques are used, making the estimator “doubly-robust”. This means that it is consistent as long as the covariate conditioning is correctly modelled by either one (or both) of the two adjustments. In this approach, we are thus not constraining our sample by matching on covariates as done in the TWFE setting, but by conditioning on the covariates by using both weighting and adjustment approaches. We thus only use information of the control variable in the baseline period (before the treatment) and avoid potential bias from the effect of the treatment on the controls themselves. We apply [Algorithm 1](#) up to step 2, such that our base sample is the same, but we do not perform the matching from step 3 onwards.⁶ The following equation specifies the estimated ATT, per group, g and treatment year, t .

$$\begin{aligned}\hat{\alpha}_{t,g} &= \frac{1}{N} \sum_{i \in \mathcal{I}} [\hat{w}_{i,g}^{treat}(y_{it} - y_{i,b}) - \hat{w}_{i,g}^{cont}(y_{it} - y_{i,b}) - \\ &\quad (\hat{w}_{i,g}^{treat} - \hat{w}_{i,g}^{cont})(\hat{m}_{i,t,g}^{cont}(X_i, \hat{\lambda}_{g,t}))] \\ \hat{\alpha}_{t,g} &= \frac{1}{N} \sum_{i \in \mathcal{I}} \underbrace{(\hat{w}_{i,g}^{treat} - \hat{w}_{i,g}^{cont})}_{\text{Inv. prob.}} (y_{it} - y_{i,b}) - \underbrace{\hat{m}_{i,t,g}^{cont}(X_i, \hat{\lambda}_{g,t})}_{\text{Outc. reg.}}\end{aligned}\tag{4.2}$$

⁶Even though this methodology significantly reduces our sample, by doing so we eliminate data inconsistencies that can arise in the CBS data and can imply that individual coefficients are driven by outliers that only appear in one or two years. We enforce that we can observe both treated and control firms for at least six consecutive years around the treatment period.

with N the amount of all firms and \mathcal{I} the set of all firms, y as one of our three dependent variables for firm, i , in the year of interest t and base year b , and X as our pre-treatment controls (whereby pre-treatment is group-specific). $\hat{w}_{i,g}^{treat} - \hat{w}_{i,g}^{cont}$ adjusts for the probability of being treated, which is dependent on all control variables and $\hat{\eta}_{i,g}^{cont}(X, \hat{\lambda}_{t,g})$ is the bias adjustment from an outcome regression. More information on both adjustments and their exact definition can be found in [Appendix B](#).

In this setting we can not enforce matching within an industry and thus use a dummy that controls for the four two-digit NACE codes that entail process-regulated firms (C17,19,23,24). Besides this, we use the same set of control variables as in the TWFE setting, but only condition on their pre-treatment levels.⁷ As in the TWFE setting we assume one year of anticipation and the base year is thus always two years before the treatment starts.

An additional advantage of this estimator is that it clearly clarifies the underlying sample of control firms. There are in principle two choices to consider, the entire population of firms that has not been treated up to t , or only the set of firms that will never be treated.⁸ For each treated group, the resulting estimator is based on weights for each set of controls that is specific for the respective treatment group. We choose to use all not yet treated firms as controls, since these will likely be similar to earlier treated firms.⁹

There is no guarantee that this set of control firms exhibits parallel trends in absence of the ETS treatment. We, however, perform a placebo test, by testing if the $\hat{\alpha}_{t,g}$ estimators for $t < g$, so before treatment starts, are significantly different from zero in a Wald test as suggested in Callaway and Sant’Anna, 2021. The results of this test will be reported with our findings.

The presented ATT in [Equation 4.2](#) allows for almost full heterogeneity, which makes it hard to draw overall conclusions. We will thus also present aggregations of the estimator to an average effect per year into the treatment, per group, and per calendar year. These aggregations are also outlined in Callaway and Sant’Anna (2021) and their exact definition can be found in [Appendix B](#), but we here discuss the underlying ideas and assumptions.

All aggregations represent a weighted sum over all available ATT estimates

⁷We have slightly adapted the notation in comparison to Callaway and Sant’Anna, 2021, since we are only considering the case of not-yet-treated firms at that moment.

⁸In the above TWFE estimation, firms that start treatment later are implicitly controls for firms that start treatment earlier, in other more standard settings, it can also happen that earlier treated firms become controls for later treated firms. However, in our above estimation even though we carefully match a control set of firms for each of the treatment groups, these controls serve as control firms for the entire set of treated firms, independent of their start in the ETS.

⁹One could argue that these firms might have expected to become treated in the future and might thus be poor controls. We will include robustness analysis for this in future versions.

such that

$$\hat{\theta} = \sum_{t=2005}^{2019} \sum_{g \in \{2005, 2008, 2013\}} \hat{w}(t, g) \hat{\alpha}_{t,g} \quad (4.3)$$

where the choice of the weights leads to various aggregations depending on the assumptions made. Estimating dynamic treatment effects, meaning the effect for all treated units e years after the beginning of their treatment, implies that we assume that the effect in 2009 is the same for a firm that starts treatment in 2005 as the effect in 2017 for a firm that starts treatment in 2013, i.e. the effect four years into treatment.

The group-specific aggregation ignores these dynamics but estimates an individual effect for each group. The advantage here is that firms starting treatment in the same phase might be more comparable, because they start with the same set of regulations and might also be treated based on similar criteria.

The calendar-year-specific estimate determines an effect per calendar year. Such specification would assume that the treatment effect is year specific, but not necessarily dynamic, since the effect in for example 2013 would be an average of the effect of phase 3 firms treated for the first time, phase 2 firms treated for 4 years and phase 1 firms treated for 7 years.

All of these aggregations have their pitfalls, but clearly reveal their underlying assumptions, while the coefficients estimated in [Equation 4.1](#) are in fact a combination of all these three aggregations.

5 Findings

5.1 TWFE

5.1.1 Matching outcomes

We manage to match 86 out of our 119 treated firms over the three ETS phases. [Figure C.1](#) in [Appendix C](#) shows for several variables how similar they are for matched firms across treatment status in 2003 before and after matching. This is the year for which we match phase 1 firms to never-treated control firms. One can see that matching makes the distributions of treated and control firms much more similar to each other. The chosen control firms are thus much more similar to their treated counterparts than the average Dutch firm. The matching comparisons for phase 2 and phase 3, for the years 2006 and 2011 respectively, can also be found in [Appendix C](#). [Table C.1](#) provides a balancing table.

5.1.2 Regression results

The TWFE results for employment can be found in [Table 1](#). The notation aligns with [Equation 4.1](#) and in columns 1 to 3 we are restricting the unit fixed effects to industry- instead of firm fixed effects. We note that the variables of interest, the DiD coefficient that tries to capture the ATT, $ETS^p \times phase^p$ are statistically insignificant throughout the specifications. Only column 3 contains one negative estimate that is significant at the 10% level. Phase 1 seems to have had a negative effect on employment, while phase 2 and phase 3 are more inconclusive. These findings are statistically not convincing, leading us to conclude that we cannot find adverse or beneficial effects from EU ETS regulation on regulated firms' employment using the TWFE method.

For value added we find similar results. For these findings we refer to [Table D.1](#). Together with the employment results we thus conclude that using a TWFE method we do not find any significant effects of the EU ETS on competitiveness.

The results for investments are presented in [Table 2](#). The coefficient of interest is statistically insignificant across specifications, making us draw the conclusion that the EU ETS has not significantly affected regulated firms' investment behavior in fixed assets.

5.2 CS2021's DiD estimator

We present here the aggregated results, the ATT per year into the treatment, per group and per calendar year. The detailed ATT estimates by group and year can be found in [Appendix D](#). [Figure 4](#) presents the three aggregations for employment as dependent variable, [Figure D.1](#) for value added, and [Figure 5](#) for the investments into fixed assets.

Before going into the results it is worth noting that we test for difference in the pre-treatment trends and can not reject similarity between the control and treatment groups. This is certainly only a proxy for the parallel trends assumption, but makes us optimistic about our approach. The overlap condition between treated and control firms is fulfilled in all periods.

The first result that first sparks attention is that for neither of the aggregations we can find any significantly negative effect of ETS participation on any of our three outcome measures. In fact we can see that the effect is significantly positive for the group that started treatment in 2005, again for all three outcome measures.

Starting with the effects on employment, in the detailed estimates of the first treated group in [Table D.2](#), one can see that the magnitude of this effect is increasing until the mid-2010s and is then slightly decreasing again. The magnitude of the estimates suggests that ETS firms that started in 2005 had a two percent higher

Table 1: *TWFE results for employment.*

	1	2	3	4	5	6
ETS ¹ × phase1	-0.013 (0.042)	-0.028 (0.053)	-0.042* (0.022)	-0.010 (0.032)	-0.014 (0.031)	-0.025 (0.019)
ETS ² × phase2	0.001 (0.056)	-0.004 (0.050)	0.00004 (0.020)	-0.010 (0.043)	0.013 (0.033)	0.011 (0.021)
ETS ³ × phase3	0.081 (0.081)	0.008 (0.057)	-0.026 (0.026)	0.066 (0.061)	0.046 (0.042)	0.002 (0.026)
phase1	0.013 (0.025)	0.031 (0.021)	0.112*** (0.015)	0.028 (0.024)	0.012 (0.018)	0.065*** (0.014)
phase2	-0.036 (0.035)	0.061** (0.028)	0.108*** (0.015)	-0.018 (0.031)	-0.005 (0.021)	0.047*** (0.014)
phase3	-0.058 (0.058)	-0.016 (0.036)	0.033* (0.017)	-0.068 (0.045)	-0.057* (0.030)	0.003 (0.019)
ETS ¹	0.432** (0.178)	-0.192** (0.089)	-0.037 (0.041)			
ETS ²	0.421** (0.163)	0.105 (0.074)	0.015 (0.032)			
ETS ³	-0.167 (0.256)	-0.314* (0.167)	-0.049 (0.056)			
Turnover		0.549*** (0.035)	-0.050** (0.024)		0.394*** (0.060)	0.010 (0.032)
Energy expenses			-0.039** (0.016)			0.0003 (0.014)
Value added			-0.023 (0.026)			0.005 (0.013)
Fixed assets investments		0.164*** (0.019)	0.051*** (0.018)		0.032*** (0.008)	0.030*** (0.010)
Wages			0.949*** (0.031)			0.561*** (0.080)
Firm FEs	No	No	No	Yes	Yes	Yes
Industry FEs	Yes	Yes	Yes	No	No	No
Observations	4,286	4,172	4,112	4,289	4,175	4,115
R ²	0.095	0.691	0.923	0.886	0.944	0.967
Adjusted R ²	0.091	0.690	0.923	0.879	0.941	0.965
Residual Std. Error	0.918	0.520	0.257	0.335	0.227	0.173

The dependent variable is the log of the number of employees. ETS^p refers to the ETS firms first regulated in phase p . The interaction $ETS^p \times phase^p$ therefore provides the DiD estimator. Control variables are in logs. Industry fixed effects are on the two-digit level. Control variables are further deflated when relevant. Standard errors are in brackets and stars refer to *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$.

Table 2: *TWFE results for investments in fixed assets.*

	1	2	3	4	5	6
ETS ¹ × phase1	0.012 (0.164)	-0.006 (0.155)	0.033 (0.155)	0.023 (0.161)	0.012 (0.156)	0.034 (0.155)
ETS ² × phase2	0.057 (0.117)	0.073 (0.110)	0.038 (0.109)	0.063 (0.115)	0.080 (0.110)	0.068 (0.111)
ETS ³ × phase3	-0.066 (0.127)	-0.106 (0.103)	-0.101 (0.100)	-0.045 (0.116)	-0.038 (0.105)	-0.026 (0.101)
phase1	-0.161*** (0.060)	-0.138** (0.056)	-0.263*** (0.056)	-0.167*** (0.059)	-0.167*** (0.057)	-0.227*** (0.055)
phase2	-0.386*** (0.079)	-0.314*** (0.066)	-0.427*** (0.069)	-0.406*** (0.077)	-0.405*** (0.072)	-0.437*** (0.071)
phase3	-0.129 (0.088)	-0.134* (0.073)	-0.152** (0.068)	-0.169** (0.085)	-0.192** (0.077)	-0.208*** (0.071)
ETS ¹	0.972*** (0.226)	0.289** (0.123)	0.036 (0.104)			
ETS ²	0.648*** (0.190)	0.215* (0.118)	0.039 (0.107)			
ETS ³	0.331 (0.212)	0.250** (0.115)	0.184* (0.102)			
Turnover		0.265*** (0.080)	0.120** (0.060)		0.296** (0.116)	0.326*** (0.105)
Energy expenses			0.251*** (0.033)			0.034 (0.048)
Employees			0.806*** (0.155)			0.897*** (0.185)
Value added		0.566*** (0.065)	0.347*** (0.067)		0.070 (0.070)	0.055 (0.071)
Wages			-0.432** (0.177)			-0.413** (0.161)
Firm FEs	No	No	No	Yes	Yes	Yes
Industry FEs	Yes	Yes	Yes	No	No	No
Observations	4,179	4,147	4,112	4,182	4,15	4,115
R ²	0.160	0.466	0.531	0.594	0.606	0.619
Adjusted R ²	0.156	0.463	0.528	0.568	0.581	0.595
Residual Std. Error	1.374	1.095	1.026	0.983	0.967	0.951

The dependent variable is the log of deflated investments in fixed assets. ETS^p refers to the ETS firms first regulated in phase p . The interaction $ETS^p \times phase^p$ therefore provides the DiD estimator. Control variables are in logs. Industry fixed effects are on the two-digit level. Control variables are further deflated when relevant. Standard errors are in brackets and stars refer to *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$.

employment than they would have had without participating in the treatment. These effects are significant at the 5 percent level.

The aggregated effects by calendar year and year into the treatment are largest in the periods in which only the first treatment group is used for the aggregations, aligning with the considerably smaller effect sizes of these groups. For those two other groups, we can not establish significantly positive or negative effects, but one can see that also for the group that started treatment in 2008, the sign is mostly positive, and highest in the mid 2010s. The sample size for the last treatment group might be too small to establish significance, but it is noteworthy that coefficient sign flips for that group.

We do not find that the higher prices of the last years have significantly reversed the positive effect, even though the 2013-group coefficient estimate is negative and the coefficients of the other two groups seem to become lower since around 2017. These might be signs that we could observe losses in competitiveness under higher prices, but until now, there is no compelling evidence for it.

When looking at the results for value added (Figure D.1), it is noteworthy how similar the results appear. The group of firms treated in 2005 is again the only one for which we can establish a significant, positive effect. It thus seems evident that the treatment has not harmed the competitiveness of any of the treatment groups.

The results for investments into fixed assets, again, look similar to the previous results, but exhibit further insights when looking closer into the dynamics of the effect. As before, only the first treatment group exhibits a significant response to the treatment, which is again positive. The magnitude is larger than before and shows that investment was 5 percent higher than without the regulation. This gives considerable support to the Porter hypothesis, suggesting that regulated firms spend a significant amount into new investments due to their regulation.

The effect is increasing considerably with the length of the exposure to the treatment for the first treatment group and one can see no reversal back to zero in the last years of high ETS prices. For the other two treatment groups, the last trading period leads to almost entirely negative coefficient estimates, except for year 2018, which for all treated groups was a positive outlier in terms of investment volumes.

5.3 Discussion

Our findings mostly point towards responses in the phase 1 group and less in the phase 2 and phase 3 groups. It is unclear if the effect for this 2005 group is so different because they are treated earlier or because they are fundamentally different. As we showed in [section 3](#) these firms were the largest firms both in terms of energy expenses and company size. Potentially this set of firms gained

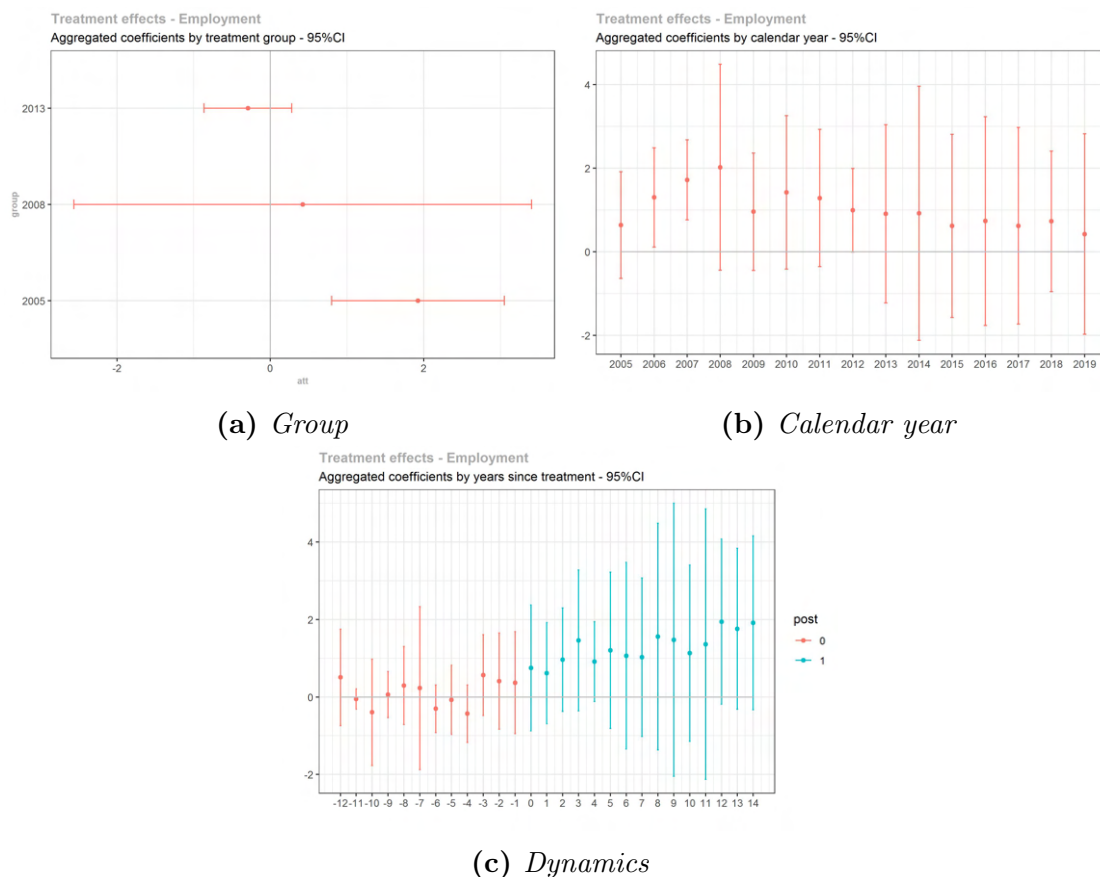


Figure 4: *Employment coefficient aggregates from analysis of Equation 4.2.*

from regulation in a Porter hypothesis-obeying fashion, by adopting to a new “greener” business model without being financially constrained by the still low prices of carbon.

Other explanations are that our selected control groups for the phase 1 firms have fundamental differences from the phase 1 firms. As we adopt a DiD design, not only the response of the treated firms matters, but also the post-treatment trends of the control groups. Inspecting Figure 3 shows that the control group also experiences changes in the variables of interest. When extending the analysis in this paper we will further dig into the representativeness of the control group for phase 1 firms.

The difference in results between the TWFE approach and the group specific ATTs can lie in the different underlying assumptions but is nevertheless surprising. The treatment coefficient for the first phase in the TWFE setting should correspond to an aggregation of the effect for the first group in the first three years of treatment in the Callaway and Sant’Anna, 2021 setting. At least for employment

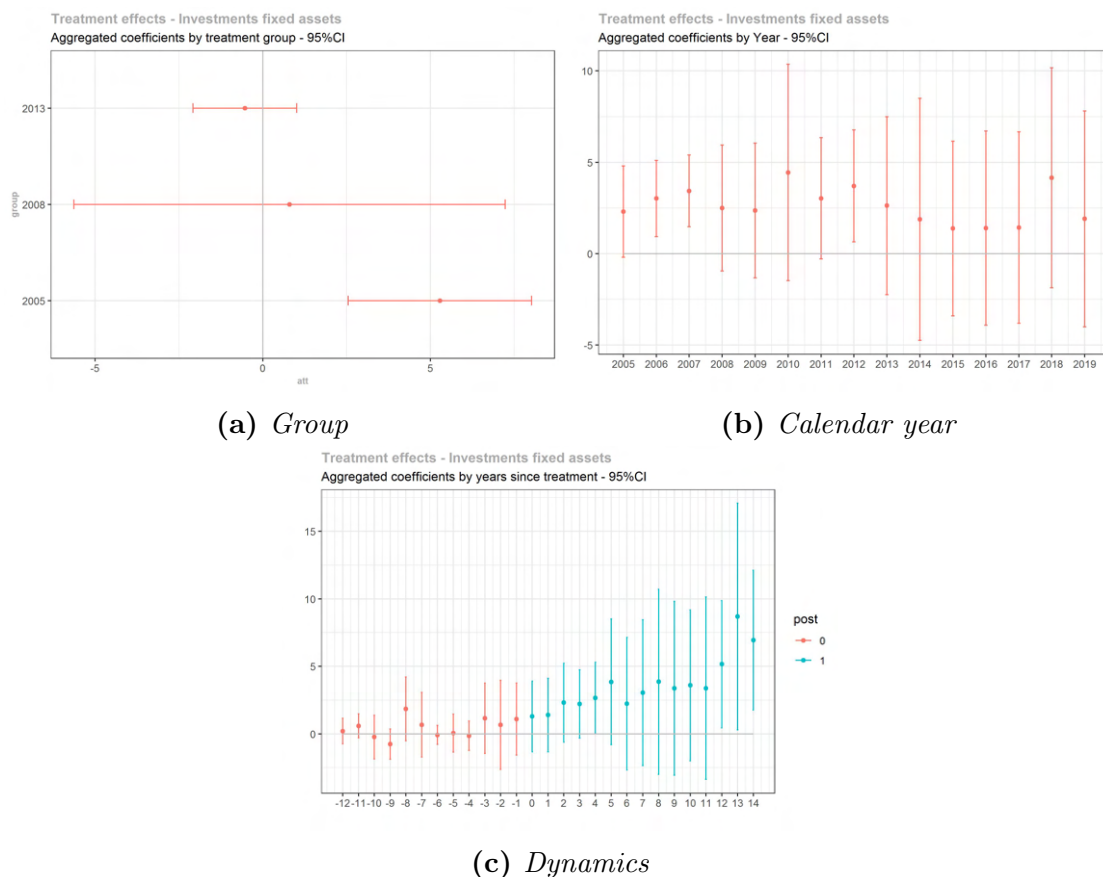


Figure 5: *Investments into fixed assets coefficient aggregates from analysis of Equation 4.2.*

this is not the case. This can of course be due to statistical nuisance, but will be further analyzed in the next steps of this analysis.

We further intend to extend the analysis over two dimensions. First, we will incorporate firm-level trade data to further study the EU ETS's effects on leakage and competitiveness. These trade data allows us to estimate whether ETS firms' exports reduced or imports from competing non-ETS firms and sectors increased. Second, we plan to disentangle the EU ETS's effects further. Currently regulation is binary within the three phases, either firms are regulated under the EU ETS or they are not, while the EU ETS does not treat all firms the same, nor is regulation the same over time. The ETS has become increasingly more stringent over time and even within the phases, e.g. with the introduction of backloading and the MSR (see Figure 2). For this purpose, we aim to define a measure of policy stringency by firm to see how the ETS affects firms that are differently affected by the regulation.

6 Conclusion

This paper studies the effects of the EU ETS on the competitiveness and the investment behavior of regulated manufacturing firms in the Netherlands. To incorporate the potential heterogeneity between different phases and firms receiving regulation in different phases, we employ two difference-in-differences (DiD) designs to estimate the average treatment effect on the treated (ATT). We employ a classical matched two-way fixed effects (TWFE) regression and a newer more flexible DiD method introduced by Callaway and Sant’Anna (2021).

Our matching results in drastically more similar treatment and control groups. The consecutive TWFE regression provides little statistically significant proof of positive or negative effects from the EU ETS on regulated firms’ competitiveness and investment behavior. If anything, phase 1 firms responded with a reduction in employment. The firms first regulated in phase 2 and phase 3 seem not to significantly respond to EU ETS regulation.

The preferred and more flexible DiD design results in insignificant estimates for firms that start treatment in phases 2 and 3, and therefore confirms the TWFE findings over this dimension. For firms that start in phase 1, however, we find that those significantly increase their employment and value added, as well as their investments in fixed assets. For employment and value added this effect peaks in the mid 2010s and seems to be reverting to zero in the later years, in which allowance prices became higher. The effect on investments is continuously increasing until the end of our sample. The considerable difference in treatment effects between firms starting in different phases can be explained by the underlying heterogeneity between the firms. Firms regulated in phase 1 were the largest and most energy-intensive firms in the sample, indicating that they might have had the biggest potential for improvements in efficiency that could have even bolstered their competitiveness. The slight contrast to our TWFE findings will be analyzed further in the next version of this paper.

From our findings we conclude that there is little evidence of competitiveness loss in the Dutch manufacturing industry. This is good news for policy makers as environmental policy does not have to go hand in hand with costly losses in employment and production. The ETS does however seem to succeed in providing some incentives to firms to update their assets. Although only large, energy-intensive firms have invested reasonably more than their non-regulated counterparts. For these firms it seems that environmental policy can stimulate innovation or adoption of new methods and equipment.

Our results fit into the literature in two ways. First, the different findings between our matched TWFE method and the more flexible DiD method highlight the importance of the right DiD design as heavily discussed in recent literature. Further, our findings are in line with the empirical EU ETS literature in that we

find some effects from the EU ETS, but not as strong as regulators might have thought upfront. We find no loss of competitiveness amongst regulated firms and we do find some positive effect on investments amongst phase 1 regulated firms. The EU ETS literature has provided several reasons for these moderate findings, amongst which low allowance prices, overallocation of (free) allowances, and two recessions that smothered production and emissions.

References

- Abadie, A. (2005). Semiparametric Difference-in-Differences Estimators. *The Review of Economic Studies*, 72(1), 1–19.
- Calel, R., & Dechezleprêtre, A. (2016). Environmental policy and directed technological change: Evidence from the european carbon market. *Review of economics and statistics*, 98(1), 173–191.
- Callaway, B., & Sant’Anna, P. H. (2021). Difference-in-differences with multiple time periods [Themed Issue: Treatment Effect 1]. *Journal of Econometrics*, 225(2), 200–230.
- Daw, J. R., & Hatfield, L. A. (2018). Matching and regression to the mean in difference-in-differences analysis. *Health services research*, 53(6), 4138–4156.
- de Chaisemartin, C., & D’Haultfoeuille, X. (2022). *Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey* (Working Paper No. 29691). National Bureau of Economic Research.
- Dechezleprêtre, A., Gennaioli, C., Martin, R., Muûls, M., & Stoerk, T. (2019). *Searching for carbon leaks in multinational companies* (Discussion Paper No. 1601). Centre for Economic Performance, LSE.
- 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.
- European Commission. (2004). Commission decision of 29 october 2004 concerning the temporary exclusion of certain installations by the netherlands from the community emissions trading scheme pursuant to article 27 of directive 2003/87/ec of the european parliament and of the council [Decision C(2004)4240-3].
- European Commission. (2005). Commission decision of 22/iii/2005 concerning the temporary exclusion of certain installations by the netherlands from the community emissions trading scheme pursuant to article 27 of directive 2003/87/ec of the european parliament and of the council [Decision C(2005) 866 final].
- European Parliament, Council of the European Union. (2003). Establishing a scheme for greenhouse gas emission allowance trading within the community [Directive 2003/87/EC].
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing [Themed Issue: Treatment Effect 1]. *Journal of Econometrics*, 225(2), 254–277.
- Heckman, J. J., Ichimura, H., & Todd, P. E. (1997). Matching As An Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme. *The Review of Economic Studies*, 64(4), 605–654.
- Jaraite-Kažukauske, J., & Di Maria, C. (2016). Did the EU ETS make a difference? an empirical assessment using Lithuanian firm-level data. *The Energy Journal*, 37(1).
- Klemetsen, M., Rosendahl, K. E., & Jakobsen, A. L. (2020). The impacts of the eu ets on norwegian plants’ environmental and economic performance. *Climate Change Economics*, 11(01).

- Löschel, A., Lutz, B. J., & Managi, S. (2019). The impacts of the eu ets on efficiency and economic performance—an empirical analyses for german manufacturing firms. *Resource and Energy Economics*, *56*, 71–95.
- Marin, G., Marino, M., & Pellegrin, C. (2018). The impact of the European Emission Trading Scheme on multiple measures of economic performance. *Environmental and Resource Economics*, *71*(2), 551–582.
- Martin, R., Muûls, M., De Preux, L. B., & Wagner, U. J. (2014). Industry compensation under relocation risk: A firm-level analysis of the eu emissions trading scheme. *American Economic Review*, *104*(8), 2482–2508.
- Porter, M. E., & Van der Linde, C. (1995). Toward a new conception of the environment-competitiveness relationship. *Journal of economic perspectives*, *9*(4), 97–118.
- Sant’Anna, P. H., & Zhao, J. (2020). Doubly robust difference-in-differences estimators. *Journal of Econometrics*, *219*(1), 101–122.
- Verde, S. F. (2020). The impact of the EU Emissions Trading System on competitiveness and carbon leakage: The econometric evidence. *Journal of Economic Surveys*, *34*(2), 320–343.
- Wagner, U., Muûls, M., Martin, R., & Colmer, J. (2014). The causal effects of the European Union Emissions Trading Scheme: Evidence from French manufacturing plants. *Fifth World Congress of Environmental and Resources Economists, Istanbul, Turkey*.
- Wagner, U., & Petrick, S. (2014). The impact of carbon trading on industry: Evidence from german manufacturing firms. (C11-V3).

A Data details

A.1 Statistics Netherlands (CBS)

The units in the CBS data are partially constructed by CBS itself. Especially the Business Unit (BE) is a construct that is generated by CBS. Here we will discuss how these units are constructed.

A.1.1 Business Unit (BE)

The business unit (BE) captures outward-facing (i.e. non-internal) Dutch production or service-provision that can be seen as one unit. This means that legal firm structures are grouped by purpose into BEs, e.g. a unit producing wooden furniture. This provides several advantages and disadvantages. The main advantage is that the BE is a unit structure that captures economic activity well. Legal firm structures often only exist for fiscal reasons and do not represent economic activity or choices well. The disadvantage is that BEs are constructed and that their composition can change over time, even though these changes might be representative of economic activity within the BE.

A.2 EU ETS

For the data on the EU ETS, coming from EUETS.INFO, a few transformations are needed.

The main problem occurs when installations change owner. This event is poorly captured by the data and therefore requires manual corrections. The corrections of ownership change were done in the following steps.

1. From the European Commission's Union Registry the lists of (stationary) installations for each phase are downloaded.¹⁰
2. The owners of each installation are compared across phases. If the owners are unchanged between phases, they are assumed to have been the same within that phase.
3. For the installations of which owners have changed between phases, we search the internet for further information to determine whether there was a transfer of ownership and between whom. From sources like news articles or websites that provide ownership data, we deduce when ownership has changed and to who. Two common situations occur, namely (1) ownership of installations is transferred within a firm group, which effectively means the installation

¹⁰These lists can be found for Phase 1, 2 and 3 on the EC's [website](http://EUETS.INFO).

has the same ultimate owner and (2) another firm purchases the installation, sometimes because the previous owner went bankrupt.

4. For installations that saw their owner change but for which we find no information when this took place, we assumed the change to take place on the day the new phase started.

The dates of ownership change then have to be reconciled with the annual data. For this, one year was chosen in which the ownership change has taken place and this year is considered to be the year in which the new owner takes economic responsibility of the installation.

A.3 Details on merging the EUTL with CBS data

Data that is imported into the CBS environment and that is identified on the chamber of commerce (in Dutch: KvK) number, like the ETS data, is encrypted on the same level. So installations under the EU ETS are imported into the CBS environment and encrypted. Encrypted chamber of commerce numbers can then be used to link EU ETS regulation to the business units.

Based on this encryption, one can find the corresponding CBS person (Dutch: persoon) in each year. This CBS person presents a layer in between the detailed KvK number and the final identifier level, which is a broader definition of firms or entity within a firm, created by CBS, namely business units (BEs). The CBS person itself is just a one to one linking from the KvK number to a CBS internal identifier. In some rare years a KvK number is assigned to two CBS persons within a year. This is because CBS draws from multiple sources which can cause duplicate links. In these cases, we have currently decided to assign the KvK number to the later created CBS person within that year.

The final identifier, the company identifier (BEID, Dutch for business unit identifier), is then a collection of one or several CBS persons that comprise one business unit within a company. The original ETS plant is thus assigned to a BEID in each year, ownership changes between years are thus uncritically represented here. However, in some years a CBS person is assigned to two BEIDs, which can happen if ownership changes within a year. In these cases, we currently assign the later BEID to the plant.

The CBS datasets are all identified on the BEID level and so we can in the next step merge the ETS plants to the CBS data sets. In each of these steps some of the companies cannot be assigned to another identifier or data set, such that in the end not all ETS firms can be merged. We don't see a reason for any systemic bias in this.

B Technicalities of estimation strategy

B.1 Further explanation and definitions of the group-year specific ATT

We here give the definitions of the inverse probability and outcome regression adjustments as well as their underlying interpretation.

$$\hat{w}_{i,g}^{treat} = \frac{G_{i,g}}{\frac{1}{N} \sum_i G_{i,g}} \quad (\text{B.1})$$

$$\hat{w}_{i,g}^{treat} = C_{i,g} \frac{\frac{p_{i,g}(X, \hat{p}_{i,g})}{1-p_{i,g}(X, \hat{p}_{i,g})}}{\frac{1}{N} \sum_i \frac{p_{i,g}(X, \hat{p}_{i,g})}{1-p_{i,g}(X, \hat{p}_{i,g})}} \quad (\text{B.2})$$

with $G_{i,g}$ being a dummy for if a firm is in the respective treatment group or not, $C_{i,g}$, a dummy that is one if the firm can serve as a control for that treatment group, thus incorporating never treated as well as not yet treated firms, and $p_{i,g}$ as the estimated propensity score for each firm (giving the probability of being in that treatment group), based on the controls and the estimated coefficients $\hat{\pi}_g$ from a logistic regression model. This procedure thus weights controls that are more likely to be treated higher than firms that are unlikely to be treated.

$\hat{m}_{i,g}^{cont}(X_i, \hat{\lambda}_{g,t})$ is the estimator of $\mathbb{E}[Y_t - Y_{base} | X, C = 1]$. It is thus the difference in predicted values between year t and the base year for the treated firms, if they were untreated. By including it, we adjust the DiD estimate by the predicted trend of the treated firms.

B.2 Further explanation and definitions on the applied aggregations

We are presenting three aggregations in this study. Once per treatment group, $g \in (2005, 2008, 2013)$, one per calendar year and one per year into the treatment. [Table B.1](#) shows the different sets of weights that these three aggregations apply. e is hereby the event time, e.g. how many years we are in the treatment, \tilde{g} is the group for which we aggregate for, and \tilde{t} the respective year of interest. Each aggregation thus takes an average over all ATT estimates that fall under the respective category and scales each ATT estimate by the total number of firms that contributed to that estimate.

Table B.1: *Weights used in the different aggregations.*

Aggregation type	$w(t, g)$
Year into treatment	$\mathbb{1}(g + e \leq 2019)\mathbb{1}(t - g = e)P(G = g G + e \leq 2019)$
Group	$\mathbb{1}(t \geq g)\mathbb{1}(g = \tilde{g})/(2019 - g - 1)$
Calendar year	$\mathbb{1}(t \leq g)\mathbb{1}(t = \tilde{t})P(G = g G \leq t)$

C Matching results


Table C.1 provides the balancing table after matching. **Figure C.1**, **Figure C.2** and **Figure C.3** show the distributions of selected variable for regulated versus non-regulated firms before and after the matching procedure for the pre-phase 1 year 2003, pre-phase 2 year 2006 and the pre-phase 3 year 2011 respectively.

Table C.1: *Balance table for the overall match outcomes for all matching years. Monetary variables are in thousands of 2015 Euros. Standard deviations are in brackets.*

	Matched Control	Treated	Difference
Value added	25,674 (29,043)	44,880 (48,839)	19,206 (6,150)
Energy expenses	2,522 (6,099)	7,843 (13,901)	5,320 (1,681)
Wages	11,568 (11,147)	19,412 (20,776)	7,843 (2,574)
No. of employees	298 (274)	436 (398)	138 (52)
Turnover	106,764 (127,009)	180,691 (198,921)	73,926 (25,373)
N	131	76	

Distributions for No. of employees

Year = 2003



(a) *Number of employees*

Distributions for Turnover (Millions EUR)

Year = 2003




(b) *Turnover*

Figure C.1: *Distributions of variables before and after matching for treated and control firms in 2003.*

Distributions for Wages (Millions EUR)


Year = 2003



(c) *Total wages*

Distributions for Energy expenses (Millions EUR)

Year = 2003




(d) *Energy expenditures*

Figure C.1: *Distributions of variables before and after matching for treated and control firms in 2003. (Cont'd.)*

Distributions for No. of employees


Year = 2006



(a) *Number of employees*

Distributions for Turnover (Millions EUR)

Year = 2006




(b) *Turnover*

Figure C.2: *Distributions of variables before and after matching for treated and control firms in 2006.*

Distributions for Wages (Millions EUR)


Year = 2006



(c) *Total wages*

Distributions for Energy expenses (Millions EUR)

Year = 2006




(d) *Energy expenditures*

Figure C.2: *Distributions of variables before and after matching for treated and control firms in 2006. (Cont'd.)*

Distributions for No. of employees


Year = 2011



(a) *Number of employees*

Distributions for Turnover (Millions EUR)

Year = 2011




(b) *Turnover*

Figure C.3: *Distributions of variables before and after matching for treated and control firms in 2011.*

Distributions for Wages (Millions EUR)

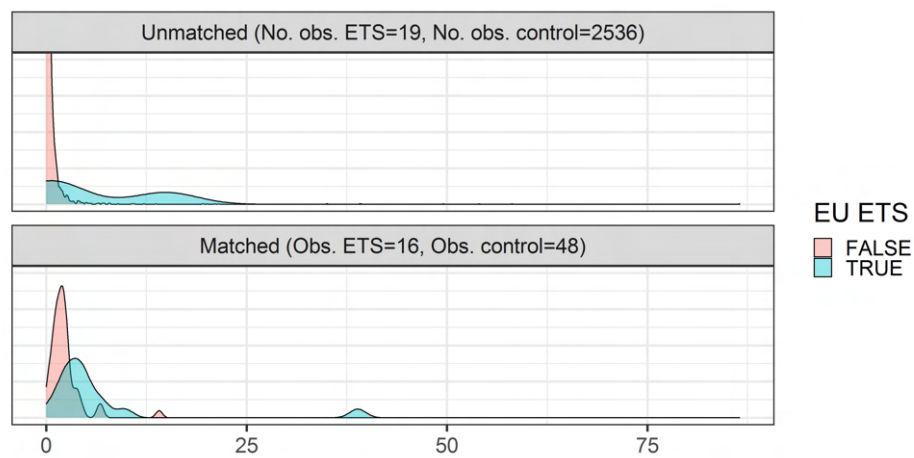
Year = 2011



(c) *Total wages*

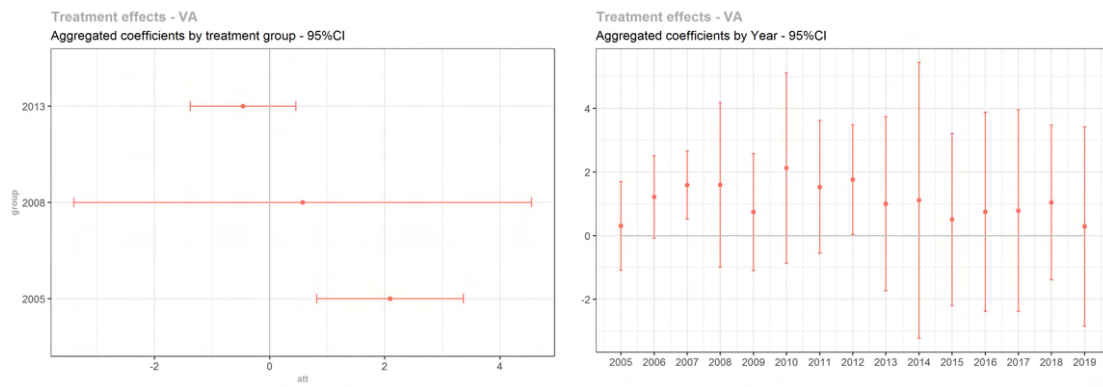
Distributions for Energy expenses (Millions EUR)

Year = 2011



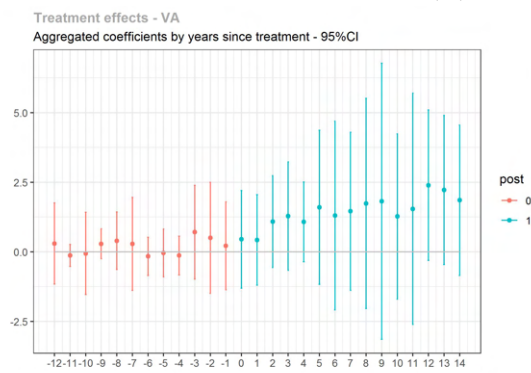
(d) *Energy expenditures*

Figure C.3: *Distributions of variables before and after matching for treated and control firms in 2011. (Cont'd.)*



(a) *Group*

(b) *Calendar year*



(c) *Dynamics*

Figure D.1: Value added coefficient aggregates from analysis of *Equation 4.2*.

D Additional tables and figures

Table D.1: *TWFE results for value added.*

	1	2	3	4	5	6
ETS ¹ × phase1	0.014 (0.066)	0.007 (0.043)	-0.001 (0.038)	0.029 (0.057)	0.013 (0.037)	0.007 (0.036)
ETS ¹ × phase2	-0.091 (0.066)	-0.068 (0.050)	-0.066 (0.041)	-0.101* (0.058)	-0.062 (0.043)	-0.066 (0.040)
ETS ¹ × phase3	0.064 (0.087)	0.009 (0.051)	-0.023 (0.040)	0.030 (0.074)	0.027 (0.047)	-0.006 (0.041)
phase1	-0.040 (0.034)	-0.044** (0.021)	0.022 (0.022)	-0.035 (0.031)	-0.065*** (0.020)	-0.015 (0.020)
phase2	-0.146*** (0.042)	-0.055** (0.028)	-0.017 (0.023)	-0.129*** (0.041)	-0.103*** (0.028)	-0.058** (0.025)
phase3	-0.068 (0.060)	-0.067** (0.032)	-0.027 (0.026)	-0.090 (0.055)	-0.116*** (0.032)	-0.077*** (0.028)
ETS ¹	0.801*** (0.193)	0.044 (0.080)	0.152** (0.066)			
ETS ²	0.608*** (0.174)	0.228*** (0.065)	0.162*** (0.060)			
ETS ³	0.060 (0.212)	-0.115 (0.114)	0.073 (0.085)			
Turnover		0.756*** (0.030)	0.310*** (0.036)		0.807*** (0.047)	0.498*** (0.057)
Energy expenses			-0.031 (0.023)			-0.025 (0.016)
Employees			-0.055 (0.056)			0.017 (0.042)
Fixed assets investments		0.122*** (0.017)	0.051*** (0.010)		0.008 (0.008)	0.006 (0.007)
Wages			0.713*** (0.063)			0.487*** (0.058)
Firm FEs	No	No	No	Yes	Yes	Yes
Industry FEs	Yes	Yes	Yes	No	No	No
Observations	4,256	4,147	4,112	4,259	4,15	4,115
R ²	0.144	0.791	0.869	0.813	0.919	0.927
Adjusted R ²	0.140	0.790	0.868	0.802	0.913	0.922
Residual Std. Error	1.039	0.507	0.395	0.499	0.326	0.304

The dependent variable is the log of deflated value added. ETS^p refers to the ETS firms first regulated in phase p . The interaction $ETS^p \times phase^p$ therefore provides the DiD estimator. Control variables are in logs. Industry fixed effects are on the two-digit level. Control variables are further deflated when relevant. Standard errors are in brackets and stars refer to *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$.

Table D.2: *Group-specific ATT estimates from Equation 4.2 for phase 1 firms.*

Year	Employment		Value added		Investments	
	ATT	SE	ATT	SE	ATT	SE
2001	-0.42	0.44	0.25	0.34	0.21	0.71
2002	0.93	0.63	1.02	0.71	1.59	1.3
2003	-0.21	0.61	-0.2	1.15	-0.42	1.79
2004	0.35	0.6	0.38	0.73	2.11	1.01
2005	0.64	0.52	0.31	0.59	2.3	1.05
2006	1.3	0.53	1.22	0.58	3.02	0.93
2007	1.72	0.42	1.59	0.48	3.43	0.86
2008	2.89	0.75	2.5	0.81	4.38	1.4
2009	1.71	0.48	1.56	0.76	4.29	1.63
2010	2.16	0.88	2.93	1.42	6.43	2.69
2011	1.98	0.63	2.24	0.89	4.95	1.65
2012	1.58	0.39	2.4	0.7	5.43	1.24
2013	2.33	0.71	2.47	0.91	6.85	1.71
2014	2.5	1.08	2.89	1.44	6.27	2.29
2015	2.12	0.69	2.16	0.88	5.44	1.72
2016	2.33	0.89	2.67	1.08	5.7	2.01
2017	1.94	0.81	2.39	1.02	5.16	1.97
2018	1.76	0.77	2.22	1.07	8.68	3.18
2019	1.91	0.86	1.85	1.07	6.93	2.02
N-treated	48		48		48	
N-Control	2988		2988		2988	

Table D.3: *Group-specific ATT estimates from Equation 4.2 for phase 2 firms.*

Year	Employment		Value added		Investments	
	ATT	SE	ATT	SE	ATT	SE
2001	-0.33	0.41	-0.26	0.46	-0.05	0.46
2002	-0.37	0.31	-0.26	0.36	-0.39	0.31
2003	-0.27	0.22	-0.25	0.22	0.05	0.41
2004	-0.01	0.17	-0.03	0.23	0.3	0.32
2005	0.25	0.81	0.47	1.08	0.63	1.47
2006	1	0.98	1.16	1.22	1.76	1.94
2007	0.55	1.24	0.22	1.5	0.38	2.34
2008	1.16	1.3	0.69	1.48	0.62	2.13
2009	0.21	1.05	-0.08	1.34	0.42	2.5
2010	0.68	1.09	1.32	1.6	2.46	2.63
2011	0.59	0.97	0.82	1.12	1.11	1.57
2012	0.41	0.8	1.12	1.08	1.98	1.82
2013	0.27	1.43	0.23	2.01	1.14	3.48
2014	0.7	2.07	1.15	3.02	0.48	4.59
2015	0.32	1.6	0.25	2.35	-0.02	3.84
2016	0.57	1.8	0.79	2.69	0.03	4.1
2017	0.16	1.7	0.42	2.58	-0.32	3.75
2018	-0.13	1.22	0.12	1.59	1.24	3.44
2019	0.12	1.54	0.09	2.43	0.41	4.41
N-treated	40		40		40	
N-Control	2988		2988		2988	

Table D.4: *Group-specific ATT estimates from Equation 4.2 for phase 3 firms.*

Year	Employment		Value added		Investments	
	ATT	SE	ATT	SE	ATT	SE
2001	0.5	0.48	0.29	0.59	0.2	0.38
2002	-0.05	0.11	-0.13	0.15	0.58	0.37
2003	-0.4	0.54	-0.06	0.56	-0.24	0.6
2004	0.06	0.24	0.29	0.24	-0.76	0.46
2005	0.3	0.38	0.39	0.38	1.84	0.96
2006	1.49	2.35	1.49	1.87	2.3	2.77
2007	-0.16	0.3	0.05	0.29	0.61	0.61
2008	0.36	1.17	0.43	1.06	0.07	1.59
2009	-1.42	1.06	-1.49	0.87	-2.15	1.37
2010	0.21	0.52	0.32	0.86	1.07	1.34
2011	0.84	0.74	1.04	1.1	1.28	1.51
2012	-0.02	0.22	-0.29	0.3	-0.22	0.63
2013	0.13	0.33	0.31	0.5	-0.11	0.84
2014	-0.45	0.37	-0.71	0.48	-1.13	0.96
2015	-0.58	0.38	-0.88	0.67	-1.29	0.83
2016	-0.69	0.61	-1.21	0.99	-1.56	1.26
2017	-0.25	0.31	-0.45	0.47	-0.54	0.74
2018	0.55	0.66	0.8	1.01	2.55	2.84
2019	-0.76	0.49	-1.07	0.77	-1.63	1.11
N-treated	17		17		17	
N-Control	2988		2988		2988	