

Themann, Michael; Koch, Nicolas

Working Paper

Catching up and falling behind: Cross-country evidence on the impact of the EU ETS on firm productivity

Ruhr Economic Papers, No. 904

Provided in Cooperation with:

RWI – Leibniz-Institut für Wirtschaftsforschung, Essen

Suggested Citation: Themann, Michael; Koch, Nicolas (2021) : Catching up and falling behind: Cross-country evidence on the impact of the EU ETS on firm productivity, Ruhr Economic Papers, No. 904, ISBN 978-3-96973-045-4, RWI - Leibniz-Institut für Wirtschaftsforschung, Essen, <https://doi.org/10.4419/96973045>

This Version is available at:

<https://hdl.handle.net/10419/233880>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Michael Themann
Nicolas Koch

**Catching up and Falling behind:
Cross-Country Evidence on the Impact
of the EU ETS on Firm Productivity**

Imprint

Ruhr Economic Papers

Published by

RWI – Leibniz-Institut für Wirtschaftsforschung
Hohenzollernstr. 1-3, 45128 Essen, Germany

Ruhr-Universität Bochum (RUB), Department of Economics
Universitätsstr. 150, 44801 Bochum, Germany

Technische Universität Dortmund, Department of Economic and Social Sciences
Vogelpothsweg 87, 44227 Dortmund, Germany

Universität Duisburg-Essen, Department of Economics
Universitätsstr. 12, 45117 Essen, Germany

Editors

Prof. Dr. Thomas K. Bauer

RUB, Department of Economics, Empirical Economics
Phone: +49 (0) 234/3 22 83 41, e-mail: thomas.bauer@rub.de

Prof. Dr. Wolfgang Leininger

Technische Universität Dortmund, Department of Economic and Social Sciences
Economics – Microeconomics
Phone: +49 (0) 231/7 55-3297, e-mail: W.Leininger@tu-dortmund.de

Prof. Dr. Volker Clausen

University of Duisburg-Essen, Department of Economics
International Economics
Phone: +49 (0) 201/1 83-3655, e-mail: vclausen@vwl.uni-due.de

Prof. Dr. Ronald Bachmann, Prof. Dr. Manuel Frondel, Prof. Dr. Torsten Schmidt,
Prof. Dr. Ansgar Wübker

RWI, Phone: +49 (0) 201/81 49-213, e-mail: presse@rwi-essen.de

Editorial Office

Sabine Weiler

RWI, Phone: +49 (0) 201/81 49-213, e-mail: sabine.weiler@rwi-essen.de

Ruhr Economic Papers #904

Responsible Editor: Manuel Frondel

All rights reserved. Essen, Germany, 2021

ISSN 1864-4872 (online) – ISBN 978-3-96973-045-4

The working papers published in the series constitute work in progress circulated to stimulate discussion and critical comments. Views expressed represent exclusively the authors' own opinions and do not necessarily reflect those of the editors.

Ruhr Economic Papers #904

Michael Themann and Nicolas Koch

**Catching up and Falling behind:
Cross-Country Evidence on the Impact
of the EU ETS on Firm Productivity**

Bibliografische Informationen der Deutschen Nationalbibliothek

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie;
detailed bibliographic data are available on the Internet at <http://dnb.dnb.de>

RWI is funded by the Federal Government and the federal state of North Rhine-Westphalia.

<http://dx.doi.org/10.4419/96973045>

ISSN 1864-4872 (online)

ISBN 978-3-96973-045-4

Michael Themann and Nicolas Koch¹

Catching up and Falling behind: Cross-Country Evidence on the Impact of the EU ETS on Firm Productivity

Abstract

This paper assesses the potential impact of the European Union Emissions Trading System (EU ETS) on firm productivity. We estimate a stylized version of the neo-Schumpeterian model, which incorporates innovation and productivity catch-up as two potential sources of firm's productivity growth, while at the same time accounting for persistent productivity dispersion within industries. This dynamic model allows us to differentiate the potential effects of the EU ETS on total factor productivity (TFP) depending on the level of firms' technological advancement. The identification approach is based on a difference-in-difference approach exploiting the incomplete participation requirements of the EU ETS and the rich panel structure of firm-level data for eight EU countries from 2002 to 2012. We find evidence that the policy effects on TFP are highly heterogeneous and depend on the distance to the technological frontier, measured as the highest TFP in each year-industry. Productivity effects are positive for firms that are close to the frontier, but they turn negative for firms operating far behind the frontier.

JEL-Code: D22, Q54, Q58

Keywords: Environmental regulation; EU ETS; productivity; competitiveness

April 2021

¹ Michael Themann, RWI; Nicolas Koch, MCC, RWI, and IZA. - We are grateful to Manuel Frondel, Christoph M. Schmidt and Nils aus dem Moore for helpful comments and suggestions. This paper also benefitted from discussions with participants at the 6th World Congress of Environmental and Resource Economists in Gothenburg and the 6th Florence School of Regulation Climate Annual Conference. - All correspondence to: Michael Themann, RWI, Hohenzollernstr. 1-3, 45128 Essen, Germany, e-mail: michael.themann@rwi-essen.de

1 Introduction

In 2005, the EU Emissions Trading System (EU ETS) was implemented as the world's largest carbon market and became the EU's flagship tool to combat climate change. This event marked a remarkable tightening of environmental regulation as firms in European energy and manufacturing industries faced a cap on their total amount of greenhouse gas emissions. Unsurprisingly, the potential economic impacts of the EU ETS have been a major issue of debate among economists, policymakers and industry representatives and play a key role in policy design since its inception.

Ambitious environmental regulation can undeniably alter decisions on production, factor allocation, investment and innovation processes and thus affect economic performance. Two broad views are prevalent in assessments of firm performance in the context of environmental policy (Ambec et al. 2013; Dechezleprêtre and Sato 2017). Conventional wisdom suggests that the policy primarily requires firms to relocate resources from traditional uses towards emissions abatement, which ultimately slows down productivity growth and diminishes their ability to compete in international markets (Jaffe et al. 1995). The second line of thought commonly referred to as the Porter Hypothesis (PH, Porter and van der Linde, Claas 1995), states that the regulation creates incentives for innovation spurring productivity growth and improving the performance of regulated companies (strong version of PH). Such gains are claimed to be more likely under market-based policies since they provide more flexibility for firms to adopt to changes in the policy environment (narrow version of PH, Jaffe and Palmer 1997).

Our study seeks to contribute to this discussion by estimating the incidence of policy-induced productivity growth stemming from the EU ETS, the grand experiment in market-based climate policy. We focus on total factor productivity (TFP), i.e. the efficiency with which firms turn inputs into outputs, as a summary estimate of the costs and benefits borne by firms (à la Greenstone et al. 2012 for the US). TFP provides us with a clear interpretation to properly assess the economic impacts of European carbon policy: it directly reflects efficiency changes due to factor reallocation (conventional wisdom) as well as effects of new technologies or innovations (Porter Hypothesis).

We apply a neo-Schumpeterian frameworks (Aghion and Howitt 2006) that characterizes TFP growth in transition to a long-run equilibrium level of TFP relative to technological leaders – the frontier (Griffith et al. 2004; Griffith et al. 2009). We model productivity growth for a follower as a positive function of the growth at the frontier (technological pass-through)

and the gap between the frontier and the productivity level of the follower (technological catch-up).

Our identification approach exploits the incomplete participation requirements of the EU ETS. It builds on a difference-in-differences approach and the rich panel structure of firm-level data for 8 EU countries from 2002 to 2012. We use the ORBIS firm database from Bureau van Dijk (BvD) and put particular emphasis on neglected challenges in constructing an international dataset that is suitable for a consistent cross-country analysis of productivity in the context of the EU ETS. First, we highlight the importance of cleaning the ORBIS data, verifying its internal consistency and correcting for any remaining biases that may affect productivity estimations. Second, we ensure that our data is internationally comparable, highly representative of each national economy in terms of coverage and size distribution and exhibits little data fluctuations. Third, we employ a novel approach for production function estimation proposed by Collard-Wexler and De Loecker (2016) that accounts for measurement error in capital stocks, a challenge prevalent in international firm-level data, and controls for unobserved productivity shocks.

The robust evidence we obtain corroborates our hypothesis that the effects of the EU ETS on TFP growth are highly heterogeneous and depend on the distance to the technological frontier, measured as the highest TFP in each year-industry. Productivity effects are positive for firms that are close to the frontier, but they turn negative for less advanced laggard firms and particularly so for firms operating far behind the frontier. This finding yields support to two central paradigms in the economics of environmental regulation, i.e. the conventional wisdom that the economic impacts of European carbon policy can entail efficiency decreases and the view that policy can induce efficiency gains as firms benefit from innovations or technology adoption from the market.

Related studies The contributions of our paper are manifold and broadly relate to two different streams of literature: First, we contribute to studies that assess the impact of environmental policies on economic performance. While evidence for the weak version of the Porter Hypothesis, that stricter environmental regulation leads to more innovation, is fairly established (Ambec et al. 2013), the evidence on the strong version of the Porter Hypothesis, that stricter regulation enhances business performance, is inconclusive. Studies on the industry level and the firm level have come up with mixed evidence in terms of significance and direction of the effect on productivity growth (Cohen and Tubb 2017; Kozluk and Zipperer 2015; Ambec et al. 2013). In a study with an approach similar to ours Albrizio

et al. (2017) use a neo-Schumpeterian productivity model and find that policy stringency increases productivity for countries and firms operating close or on the technology frontier. The effect becomes smaller in size and eventually insignificant or even negative for laggard performers.

One explanation for the mixed results is grounded in important econometric and data issues that we attempt to resolve with our research design (Becker 2011; Gray and Shadbegian 2001): We adequately control for heterogeneous firm behavior as well as unobserved industrial trends and country-level heterogeneity – a feat most prior studies lack. Cross-country studies have also suffered from a lack of cross-country policies and have been confined to utilize broad indicators of environmental stringency (Albrizio et al. 2017, compliance costs (Gray 1987; Morgenstern et al. 2002; Gray and Shadbegian 2001) or pollution intensities (Cole and Elliott 2003; Harris et al. 2002; McConnell and Schwab 1990) which raises concerns of endogeneity due to simultaneity, reverse causality or measurement error. We instead use the variation in firm regulatory status created by the EU ETS to help obtain an estimate of its potential effect on productivity growth (similar to Greenstone et al. (2012) for air quality regulation in US manufacturing). Our study focuses on the narrow version of the PH (productivity gains more likely under market-based policies) for which the empirical evidence so far is particularly scarce.

Second, we contribute to the emerging literature on the economic impacts of the EU ETS that so far has delivered mixed results for a broad range of performance indicators, such as employment, profits, revenues, output, or investments (Abrell et al. 2011; Chan et al. 2013; Commins et al. 2011; Petrick and Wagner 2014). The few papers that looked at TFP found insignificant and sometimes either marginally negative or positive effects (Marin et al. 2018; Commins et al. 2011; Löschel et al. 2019; Lutz 2016). We suggest that part of the reason for the lack of consistent empirical results may be due to incomplete consideration of the dynamic adjustments depicted in neo-Schumpeterian frameworks that are well-established in general productivity literature (Griffith et al. 2004; Griffith et al. 2009) as well as the wide range of samples and outcomes used in these studies. Our study thus contributes to existing evaluations of the EU ETS not only by providing a very comprehensive cross-country analysis with TFP as a complete measure of total costs and benefits, but by accounting for dynamic adjustments and by addressing a number of data related biases as a crucial prerequisite for consistent cross-country productivity estimations (Akerberg et al. 2015; Collard-Wexler and De Loecker 2016; Gopinath et al. 2017).

2 Empirical framework

Our empirical strategy is based on a stylized version of the Neo-Schumpeterian model of total factor productivity growth that is grounded in the theoretical models of endogenous growth (Aghion and Howitt 2006). The key idea here is to define TFP growth in transition to a long-run equilibrium level of TFP relative to the technology frontier. We follow Griffith et al. (2004) and Griffith et al. (2009)) and model productivity growth for firms lagging behind the technological leader as a positive function of growth at the European technological frontier (technological pass-through) and the gap between the frontier and the productivity level of the follower (technological catch-up).

We also account for persistent productivity dispersion within industries, an important stylized fact in the productivity literature (as in Griffith et al. (2009)). Driving forces behind this persistency are differences in the access to managerial capacity and inputs, such as labor, capital and activities related to research and development (Syverson 2011). Technologically more advanced firms are also more likely to benefit from innovations and to adopt new technologies and processes available in the market, whereas the opposite is true for technologically less advanced firms (Griffith et al. 2004). Our econometric model consequently introduces a long-run relationship in which follower firms lie a steady-state distance behind the frontier such that their rate of productivity growth including catch-up equals productivity growth at the frontier.

We adopt the view established by Bourlès et al. (2012) that a policy cannot only directly drive or curb productivity growth, but that a firm's response to the new regulatory environment may depend on their level of technological advancement. We assume firms close to the technological frontier to have the means to perform well or even thrive under a carbon policy, e.g. by having access to high quality sources of labor and capital, the right management structure or the research leadership necessary to bear not only the costs of abatement measures but to reap the benefits from green innovation. Less advanced firms that are lagging far behind the frontier, in contrast, may suffer from efficiency decreases stemming from abatement costs due to their lower ability to adopt efficient technologies and processes available in the market.

Formally, productivity A for firm-country-industry pair ics can be modeled as an autoregressive distributed lag (ADL 1,1) process that is co-integrated with the productivity at the

technology frontier F (see e.g. Griffith et al. (2004) for more details). Assuming long-run homogeneity, we can capture this process in the following form:

$$\ln A_{icst} = \alpha_1 \Delta \ln A_{Fst} + (1 - \alpha_0) \ln \left(\frac{A_{Fst}}{A_{ics}} \right)_{t-1} + \epsilon_{icst} \quad (1)$$

The first term is the growth of TFP at the frontier and depicts the technological pass-through.

A straightforward way of defining the global productivity frontier would be to take the top 5% of firms in terms of productivity levels for each industry and year (Albrizio et al. 2017). However, it is important to account for the fact that the number of firms in ORBIS grows over time, which is why we define the European (subsequently “global”) productivity frontier based on a fixed number of firms (Andrews et al. 2015): We identify the average of the number of firms that constitute the top 5% of firms in each sector-year productivity distribution for the period 2002-2012. Then, for each sector-year the global frontier is defined as the TFP of the last firm in this top group, i.e. that is operating right at the frontier. The second term captures the technological catch-up process and is defined as the gap to the country-industry frontier, i.e. the difference of TFP at the frontier and at the firm productivity level in $t-1$. α_1 thus denotes the degree to which a one percent increase in TFP growth at the frontier contributes to TFP growth at the level of the individual firm. α_0 captures by how much percent an increase in individual firm productivity level in $t-1$ increases the firm TFP level in t . $(1 - \alpha_0)$ captures the contribution of the technology gap in $t-1$ to current firm productivity growth. Productivity catch-up thus occurs if the coefficient for frontier TFP growth is $\alpha_1 > 0$ and the coefficient for the distance to the frontier is $(1 - \alpha_0) > 0$.

In line with Bourlès et al. (2012), we adapt Equation 1 for our specific policy context in order to obtain an estimate of the potential effect of the EU ETS on productivity growth. More specifically, we use the variation in firm regulatory status (measured by the variable REG_{it} , which is a dummy indicating if firm i is regulated by EU ETS at time t) to identify the average effect of the EU ETS on regulated firms’ productivity growth using a differences-in-differences framework. We allow the EU ETS regulation to have a nonlinear effect by crossing the REG_{it} variable with distance to frontier. We also include vectors of firm control variables x_{icst-1} (tangible fixed assets, operating revenue, age, multinational status) and

fixed effects d to account for observed and unobserved heterogeneity. The resulting specification is as follows:

$$\ln A_{icst} = \alpha_1 \Delta \ln A_{Fst} + (1 - \alpha_0) \ln \left(\frac{A_{Fs}}{A_{ics}} \right)_{t-1} + \beta_1 REG_{it} + \beta_2 REG_{it} \ln \left(\frac{A_{Fs}}{A_{ics}} \right)_{t-1} + \delta x_{icst-1} + \gamma d + \epsilon_{icst} \quad (2)$$

While a randomized experiment would automatically balance out important confounding factors between regulated and not regulated firms and allow a simple comparison of means to yield the average effect on the treated (ATET), this is not the case in our study. Given that we are working with observational data, covariate balance between the two groups is not necessarily the case.

Although treatment assignment under the EU ETS is not random, we can exploit the particular regulatory features of the EU ETS. To keep implementation costs low, the system only comprises firms that own large installations in carbon-intensive industries. Regulatory status is set via industry specific criteria such as capacity thresholds. For instance, a steel plant will only be regulated by the EU ETS if its hourly production capacity lies above 2.5 tons. Importantly, productivity growth can likely be explained by factors and decisions taken at the firm level (e.g. on technology, asset and employment structure, firm size) rather than at the level of a specific plant. This insight introduced by Calel and Dechezleprêtre (2016) should allow us to find a suitable group of EU ETS and control firms that fall under different regulatory regimes but are very similar in a number of important productivity determinants except for installation size. At least in principle, reasonably similar firms should then be comparable (Fowlie et al. 2012). This means that, under certain conditions, we can establish balance in a number of important confounders and thus aim to recuperate the conditions of a randomized experiment. In order to obtain a meaningful estimate of the potential regulatory impact under this setup, we thus employ entropy balancing combined with a differences-in-differences estimator.

The new third term in Equation 2 is a differences-in-differences estimator of the regulatory impact on productivity growth. It captures variation in TFP growth specific to EU ETS firms relative to non-EU ETS firms in the years after the policy was introduced (2005-2012) relative to before (2002-2004). The fourth term measures the heterogeneous effect of the policy that depends on the distance to the technology frontier in $t - 1$. The total effect of the EU ETS on firm productivity growth is, thus, given by $\beta_1 + \beta_2 \ln \left(\frac{A_{Fs}}{A_{ics}} \right)_{t-1}$.

Entropy balancing can be considered as a generalization of the propensity score approach (Hainmueller 2012). Rather than estimating each firm's propensity to be treated and thus condensing covariate information into a single score, entropy balancing enables the researcher to exploit knowledge on the covariate distribution moments. In essence, the algorithm uses a maximum entropy reweighting scheme to assign each control observation a weight such that the covariate moments of both groups are exactly identical, i.e. perfectly balanced (Hainmueller 2012; Hainmueller and Xu 2013). This approach has several advantages over traditional balancing approaches. In particular, it provides considerably improved covariate balance while relieving the researcher from running several iterations until finding a specification that minimizes differences in distribution moments.

We take full advantage of these possibilities and enforce a perfect balance in all three moments (mean, variance and skewness) for the full set of covariates that we deem potentially important confounding factors: pre-policy values for productivity growth, operating revenue, tangible fixed assets and employment for the years 2002-2004. Hence, we augment Equation 2 with a set of balancing weights obtained from this process.

The panel structure of our data enables us to account for any remaining heterogeneity that is not removed by the DiD term and the balancing process and thus may still confound our estimates. We fully exploit these benefits and account for country-industry-specific characteristics (such as technology or skills), country-specific trends (such as overall technical progress or deregulation waves), and year-specific shocks that may potentially drive any differences in TFP growth between EU ETS firms and non-ETS firms.

We also estimate specifications with a firm-specific term, which reflects heterogeneity in innovative capabilities, in our dynamic framework. While controlling for time-constant firm characteristics may be important in our policy context, the inclusion of both a lagged dependent variable and fixed effects induces a downward bias in the estimated coefficient on the lagged dependent variable (as first pointed out by Nickell (1981)). This creates an upward bias in the term that measures the lagged gap to the frontier, which contains information of firm TFP in $t - 1$.

We therefore explore the robustness of our inferences by estimating a specification with firm fixed effects using the Arellano-Bond estimator (also called difference GMM, Arellano and Bond (1991)). This estimator uses the levels of additional lags of the dependent variable as instruments for the differenced lagged dependent variable to address the simultaneity bias.

Our main identification assumption is that no unobserved variables exist that simultaneously influence changes in productivity and the probability of being regulated by the EU ETS (unconfoundedness). In our specific DiD framework it requires that in the absence of treatment the productivity of ETS and non-ETS firms must follow the same trend, the so called common trend assumption (Lechner 2010).

We also assume the absence of spillover effects occurring between the treated and the control group. This is known as the stable unit treatment value assumption (SUTVA). The third assumption is the absence of anticipation effects. Anticipating the start of the European carbon market pre-2005, firms may have had some incentive to either avoid or select themselves into the regulatory policy, e.g. by down- or upsizing operations. In rebuttal, recent analyses indicate that this incentive may have been limited or inexistent, e.g. due to the uncertainty surrounding the negotiation process on policy implementation pre-2005 (austem Moore et al. 2019). We investigate the plausibility of our assumptions in the robustness section.

Standard errors are clustered on the four-digit industry level and thus explicitly allow the error term to be correlated across time within and across firms within each four-digit sector (as in Griffith et al. (2009)). This important feature enables us to account for input–output linkages between firms and sub-industries within the same four-digit industry cluster. As demonstrated in Bertrand et al. (2004)), clustering performs very well in allowing for correlated errors in settings with at least 50 clusters, which is the case for our data.

3 Data

We construct a novel database that is suitable for measuring TFP in a cross-country setting and assessing the possible impact of the EU ETS. For this purpose, we bring together two sources of raw data: Firm financial data from ORBIS and regulatory data on the EU ETS from the European Union Transaction Log (EUTL).

3.1 Firm financial data

Our primary datasource is the ORBIS database compiled by the commercial data provider Bureau van Dijk (BvD). ORBIS collects information from administrative sources, in particular detailed balance sheets, income statements, and profit and loss accounts of firms. The financial accounting data is harmonized across countries and, along with information of

firm ownership, provided in a global standard format. The database thus allows for cross-country comparisons and is constantly being updated. BvD extracted the financial data we use in the last week of November 2015. The data contains all firms above a turnover of one million Euro, total assets of 2 million Euros or a total number of 15 employees in 2015, which amounts to a sample of around 12.5 million firms. Our financial data is reported in thousands of Euros for the years 2002-2012. We use unconsolidated financial information from local registry filings to ensure a high quality of the raw data. Industries are defined by their four-digit industry NACE Rev. 2 classification.

Our methodology to estimate TFP requires information on intermediate inputs, a variable for which ORBIS sometimes reports missing values, particularly for Great Britain.¹ To improve the coverage of the variable, we obtain imputed intermediate inputs by taking the difference between operating revenue and value added. This leads to a notable improvement in coverage for Great Britain. TFP distributions using internally imputed variables have been shown to be almost identical to their non-imputed counterparts (Gal 2013).

We then follow a thorough process of data cleaning. Our four-step procedure builds on Gopinath et al. (2017). First, we correct for reporting mistakes by e.g. dropping observations with missing information or implausible values. Second, we verify the internal consistency of balance sheet information. We construct ratios that compare the sum of variables belonging to some aggregate to their respective aggregate. For instance, we calculate the sum of tangible fixed assets, intangible fixed assets, and other fixed assets as a ratio of their respective aggregate, i.e. total fixed assets.² We estimate the distribution for each of these ratios and remove extreme values at the tails of the respective distributions (below the 0.1 percentile and above the 99.9 percentile). These first two steps are implemented at the level of the total economy for each country separately. Third, we do a more specific quality control of all variables that are part of our analysis, including our estimated measures of TFP. Lastly, we winsorize the variables at the 1 and 99 percentile.

All nominal variables used in our analysis are deflated with a corresponding yearly price deflator at the two-digit industry level. Included are price deflators on value added, gross output, intermediate inputs as well as on capital and investment goods. This procedure ensures that the growth rates of the variables used for productivity calculations are not driven

¹We use material costs as a proxy for total intermediate inputs, since energy usage and purchased services are not reported in ORBIS. While this inhibits estimating more complex output based production functions, it is still very suitable for the standard value added based function.

²This is also conducted for the following aggregates: total assets, total current assets, total shareholder funds and liabilities.

by price changes. All price deflators are retrieved from the Structural Analysis Database of the OECD.³ In case 2-digit deflators are not available we use the information from higher levels of industry aggregation. In the rare case that no specific price deflator is available for a given industry, we use value added deflators.

After deflating all financial covariates over time, we correct for price differences across countries by converting them into a common currency (USD). We fix the exchange rate at the middle of the sample period, in 2007, to mitigate the influence of fluctuating exchange rates on the productivity numbers. This is particularly suitable for industries operating under strong international competition such as manufacturing (Gal 2013).

We restrict our sample in two steps. First, we focus on the manufacturing and energy sector, since these sectors correspond to over 90% of all emissions under the EU ETS in 2012. Second, we apply the criteria suggested by Kalemli-Ozcan et al. (2015) to construct data that is representative of each European economy along three dimensions: (i) how much of the gross output from Eurostat Structural Business Statistics (Eurostat SBS)⁴ we cover from 2002-2012; (ii) the degree of data fluctuations measured by the standard deviation and (iii) the firm size distribution. We then limit our sample to the 8 countries that meet our criteria, i.e. have high coverage of the economic activity (at least 65% on average), have little fluctuations (standard deviation below 0.1) and are representative of different firm sizes (Orbis-Eurostat %-difference in each size category below 10).⁵ We consider such an approach crucial for obtaining empirically meaningful results both in terms of TFP estimations as well as our subsequent analysis with respect to the EU ETS.

3.2 Production function estimation

We estimate a Cobb-Douglas production function separately for each two-digit industry s in our sample of European firms. Firm i at time t produces value added according to the following log linear form:

$$\ln y_{ist} = \beta^{l(s)} \ln l_{ist} + \beta^{k(s)} \ln k_{ist} + \ln Z_{ist} + \epsilon_{ist}, \quad (3)$$

where y_{ist} is value added and $\ln Z_{ist}$ is the unobservable Hicks neutral productivity term. l_{ist} denotes the wage bill, and k_{ist} denotes the capital stock. $\beta^{l(s)}$ is the elasticity of value added

³The data can be retrieved from <http://www.oecd.org/sti/ind/stanstructuralanalysisdatabase.htm>.

⁴The data can be retrieved from https://ec.europa.eu/eurostat/statistics-explained/index.php/Structural_business_statistics_overview.

⁵The selected countries are Belgium, France, Germany, Great Britain, Italy, Spain, Sweden and Norway.

with respect to labor and $\beta^{k(s)}$ is the elasticity of value added with respect to capital. The elasticities are constants determined by the available technology and the same for a group of firms within a given industry s . It is a general best practice to assume homogeneity in technology at the 2-digit industry-level. The error term ϵ_{ist} captures random shocks and measurement error and we assume it to be identically and independently distributed. The capital stock k_{ist} is calculated according to the Perpetual Inventory Method (Gal 2013). It defines the level of real capital stock k_{ist} in firm i and sector s and in year t as $k_{ist} = k_{ist-1}(1 - \delta_{ist}) + i_{ist}$, where i_{ist} are real investments and δ_{ist} is the depreciation rate.

We employ the methodology by Collard-Wexler and De Loecker (2016) and refer to their paper for details on the estimation procedure. Given our estimated elasticities $\beta^{l(s)}$ and $\beta^{k(s)}$, we then calculate firm (log) productivity as $\ln Z_{ist} = \ln y_{ist} - \beta^{l(s)} \ln l_{ist} - \beta^{k(s)} \ln k_{ist}$. Note that, similar to most empirical studies, we do not observe prices at the firm-level, but at the industry-level. Our measure of TFP is thus revenue-based productivity and captures a combination of market power and productivity (Syverson 2011).

The extensive literature on production function estimation has shown that applying OLS will most likely yield biased coefficients (Ackerberg et al. 2015, ACF). The reason for this major obstacle in obtaining TFP is that the firm chooses the observed input factors that enter the production function, labor and capital. For most real world applications, it is very likely that some determinants of production exist that impact the choice of inputs but are exclusively observed by the firm. If this is the case, the OLS estimates will be biased. Several important contributions have exploited control functions to address this issue of endogenous inputs (Olley and Pakes 1996, Levinsohn and Petrin 2003, Wooldridge 2009, Ackerberg et al. 2015). We refer to Ackerberg et al. (2015) for a technical description of the endogeneity problem, the long history of this issue in production function estimation, and the main different techniques to solving it.

Collard-Wexler and De Loecker (2016) (CWL) demonstrate that commonly used approaches perform poorly in the presence of notable amounts of measurement error in capital - a phenomenon common even in firm-level data from high quality sources. This measurement error may stem from "... the difficulty to appropriately measure depreciations over a long period of time across heterogeneous assets and production processes...", which can be compounded with other reporting related factors (Collard-Wexler and De Loecker 2016). Relying on Monte-Carlo simulations and two standard firm panel data sets they find that standard techniques yield downward biased capital coefficients in such a context, which

makes capital intensive firms appear more productive than they actually are. Their findings are coherent with the general observation in the empirical literature on firm productivity that capital coefficients are often very low and sometimes even negative (Becker et al. 2006).

We apply these insights to our ORBIS data and provide a comparison of input elasticities for each sector using a range of different estimators. Apart from standard procedures, we apply a two-stage IV estimator proposed by CWL that instruments capital with lagged investments to control for the measurement error and unobserved productivity shocks (the latter a la ACF, using intermediate inputs to substitute for the shock). Note that the validity of the instrument is given by construction, since lagged investments is one element to measure capital stock.

Table 1: Capital and labor elasticities, by sector

NACE 2-digit code	OLS		Fixed effects		IV (Olley-Pakes, OP)		ACF, 1 Step		ACF, 2 Step		CWL, 2 Step	
	(1)		(2)		(3)		(4)		(5)		(6)	
	β_{l_ols}	β_{k_ols}	β_{l_FE}	β_{k_FE}	β_{l_IV}	β_{k_IV}	β_{l_F}	β_{k_1stF}	β_{l_F}	β_{k_2stF}	β_{l_F}	$\beta_{k_IV2\sim F}$
10	0.85	0.14	0.79	0.05	0.81	0.17	0.69	0.07	0.69	0.17	0.69	0.28
11	0.85	0.14	0.74	0.07	0.81	0.19	0.54	0.12	0.54	0.25	0.54	0.44
12	0.92	0.08	0.77	0.08	0.94	0.08	0.59	0.18	0.59	0.40	0.59	0.61
13	0.90	0.07	0.81	0.05	0.86	0.11	0.72	0.05	0.72	0.10	0.72	0.22
14	0.91	0.08	0.78	0.04	0.88	0.11	0.74	0.06	0.74	0.09	0.74	0.22
15	0.88	0.09	0.78	0.04	0.84	0.14	0.73	0.05	0.73	0.11	0.73	0.20
16	0.88	0.10	0.82	0.04	0.86	0.12	0.73	0.06	0.73	0.11	0.73	0.06
17	0.88	0.11	0.79	0.05	0.85	0.14	0.69	0.06	0.69	0.15	0.69	0.27
18	0.89	0.09	0.80	0.04	0.85	0.13	0.75	0.05	0.75	0.09	0.75	0.08
19	0.87	0.13	0.92	0.03	0.81	0.19	0.55	0.08	0.55	0.34	0.55	0.41
20	0.89	0.09	0.81	0.04	0.84	0.14	0.65	0.08	0.65	0.17	0.65	0.32
21	0.94	0.04	0.85	0.03	0.86	0.12	0.63	0.04	0.63	0.13	0.63	0.37
22	0.88	0.10	0.81	0.04	0.84	0.13	0.70	0.07	0.70	0.15	0.70	0.26
23	0.85	0.12	0.83	0.03	0.81	0.16	0.66	0.07	0.66	0.17	0.66	0.30
24	0.87	0.12	0.84	0.03	0.84	0.15	0.70	0.06	0.70	0.20	0.70	0.27
25	0.89	0.09	0.87	0.03	0.86	0.12	0.75	0.06	0.75	0.12	0.75	0.20
26	0.92	0.06	0.89	0.03	0.88	0.10	0.73	0.07	0.73	0.11	0.73	0.06
27	0.91	0.08	0.86	0.03	0.87	0.12	0.73	0.07	0.73	0.11	0.73	0.06
28	0.91	0.07	0.85	0.03	0.88	0.11	0.74	0.06	0.74	0.11	0.74	0.22
29	0.91	0.08	0.86	0.03	0.87	0.12	0.75	0.06	0.75	0.12	0.75	0.09
30	0.90	0.08	0.89	0.03	0.88	0.11	0.74	0.05	0.74	0.16	0.74	0.22
31	0.92	0.07	0.86	0.03	0.89	0.10	0.78	0.04	0.78	0.07	0.78	0.19
32	0.90	0.07	0.82	0.04	0.87	0.11	0.73	0.06	0.73	0.10	0.73	0.06
33	0.91	0.07	0.87	0.03	0.89	0.09	0.81	0.05	0.81	0.09	0.81	0.04
35	0.85	0.12	0.75	0.10	0.78	0.19	0.40	0.11	0.40	0.18	0.40	0.73

As shown in Table 1, our CWL estimates yield consistent capital and labor elasticities for all 25 industries. We have no cases of a zero, negative, missing or excessively big values, in which case we would omit the industry. Even more importantly, the coefficients are highly consistent when compared to alternative procedures. Capital coefficients are notably higher and, in most cases, the sum of capital and labor coefficients is close to one. In contrast, columns 1-5 demonstrate that capital coefficients obtained from standard estimators are very low which suggests the presence of measurement error and a strong bias towards zero with respect to the true parameter (Collard-Wexler and De Loecker 2016). This gives us confidence that our estimator is the most suitable for obtaining productivity in ORBIS. We

rely on the standard assumption that labor is a static input choice and exploit the first-order condition to directly estimate the labor coefficient as the input cost share of labor.

Table 2: Summary statistics

Variable	Mean	Std. Dev.	Min	p25	p50	p75	Max
TFP growth (in logs)	-0.01	0.19	-0.64	-0.09	0.00	0.09	0.58
Gap to TFP frontier (t-1)	0.70	0.41	0.00	0.45	0.68	0.91	2.73
TFP frontier growth (in logs)	0.00	0.06	-0.31	-0.03	0.00	0.03	0.58
Tangible fixed assets (in logs, t-1)	2.29	2.26	-2.57	0.73	2.11	3.60	9.58
Operating revenue (in logs, t-1)	4.43	1.85	1.38	3.08	3.95	5.31	10.98
MNE (1=yes)	0.08	0.27	0.00	0.00	0.00	0.00	1.00
Firm age	22.99	16.48	3.00	12.00	19.00	29.00	91.00

Table 2 reports summary statistics for the variables that we include in our subsequent analysis of the potential impacts of the EU ETS. It shows that, on average, the rate of TFP growth is similar to growth at the frontier (mean -0.1 and mean 0). However, growth in the full sample is highly dispersed which suggests that some firms are converging towards the frontier whereas others are falling further behind.

Table 3: TFP relative to frontier firms by country

Country ISO Code	Mean	Std. Dev.	p1	p10	p25	p50	p75	p90	N
BE	0.60	0.29	0.18	0.34	0.43	0.53	0.68	0.92	49,426
DE	0.59	0.29	0.07	0.30	0.44	0.55	0.70	0.93	94,213
ES	0.51	0.23	0.14	0.30	0.37	0.47	0.59	0.76	289,492
FR	0.58	0.23	0.22	0.35	0.44	0.54	0.67	0.84	270,961
GB	0.62	0.29	0.19	0.36	0.44	0.55	0.72	0.99	61,306
IT	0.54	0.25	0.17	0.31	0.39	0.50	0.63	0.82	455,657
NO	0.58	0.28	0.07	0.32	0.42	0.54	0.69	0.91	28,277
SE	0.55	0.26	0.07	0.31	0.40	0.51	0.64	0.84	53,542
Total									1,302,874

In Table 3 we discuss the distribution of TFP relative to the productivity frontier for each of the selected countries. The table highlights the considerable country heterogeneity that exists with respect to the size of the gap to the common European frontier. For instance, the TFP of the average firm in Great Britain constitutes 62% (Belgium: 60%) of European frontier TFP. The average British firm operates comparatively close to the frontier. In contrast, the average firm in Spain has an estimated productivity of merely 51% of the European frontier (Italy: 54%). Table 3 also lends support to the view that there is a notable degree of dispersion in our cross-country sample. It appears that the distribution of TFP has longer tails in Southern European countries. At the 90th percentile, the gap is still 24% in Spain (Italy: 20%), whereas in Great Britain it is only 1% suggesting that these firms are already

operating very closely at the frontier (Germany: 7%). While one has to consider notable differences in firm coverage among countries, taking the insights from Table 2 and 3 together shows that productivity levels and dynamics can have substantial variation (Bartelsman and Doms 2000; Syverson 2011).

3.3 Regulatory data

The EUTL is the official registry of all plants under the EU ETS and is composed of regulatory information. We utilize this plant-level data to identify firms in ORBIS that are regulated by the EU ETS. Every installation in the EUTL is owned by a so-called account holder and we can match the information provided on this latter entity (national firm identifier, name, address) with data contained in ORBIS.⁶ We utilize the national identification number that uniquely identifies a firm in both datasets. We correct for systemic errors in the EUTL to make these identifiers compatible with the harmonized format used in ORBIS. Only in very few cases we could not identify firms via this method and instead located them via their name. We successfully match 8.218 out of all 8.578 companies (96%) that as of March 2014 hold a regulated plant, which corresponds to 14.507 out of a total of 15.043 installations under the EU ETS (96%). We then test the internal consistency of our procedure and compare the companies' contact information between the EUTL and ORBIS. For 98.2 % of the identified companies, the information between both sources is highly consistent, whereas for 1.8% of firms the seemingly inconclusive information is mostly related to changes in company names or mergers and acquisitions. The remainder of not-matched entries could either not be found in ORBIS or information was incomplete, which in many cases are hospitals, governmental agencies or universities. We use the emission data from the EUTL as an indicator of activity and keep 7.279 firms that were active in phases I or II of the EU ETS.

4 Empirical results

4.1 Main results

We report balancing statistics in Table 4 for our full sample both before and after applying entropy balancing weights to the two groups of EU ETS and non-EU ETS firms. For

⁶The EUTL subsection "List of Stationary Installations in the Union Registry" contains all installations under the EU ETS as of February 27, 2014. It can be retrieved from <http://ec.europa.eu/clima/policies/ets/registry/documentation.en.htm>.

both groups, it presents the respective covariate's mean, variance and skewness. In addition, it reports standardized differences between the two distributions. The standardized difference is considered a particularly reliable measure of covariate balance as it is robust to changes in sample size and comparable across different covariate scales. Maximum values for standardized differences are recommended to be from 10 to 25 percent (Garrido et al. 2014).

Table 4: Covariate balance, entropy balancing

<i>Pre weighting</i>							
Variable (in logs)	Mean		Variance		Skewness		Stand. Diff.
	Treated	Control	Treated	Control	Treated	Control	
TFP growth, 2004	0.03	0.004	0.05	0.03	0.01	-0.21	12.54
TFP growth, 2003	-0.02	-0.005	0.04	0.03	0.05	-0.21	-6.82
Tangible fixed assets, 2004	6.16	2.68	3.91	5.35	-0.17	0.56	175.88
Tangible fixed assets, 2003	6.14	2.65	3.96	5.34	-0.19	0.55	175.41
Tangible fixed assets, 2002	6.11	2.60	3.97	5.23	-0.22	0.52	176.29
Operating revenue, 2004	7.31	4.84	4.02	4.05	0.04	1.02	122.90
Operating revenue, 2003	7.26	4.79	4.07	4.00	0.02	1.01	122.27
Operating revenue, 2002	7.23	4.73	4.10	3.90	0.00	0.98	123.22
Employment, 2004	5.21	3.33	1.95	1.40	-0.46	0.49	135.06
Employment, 2003	5.23	3.35	1.96	1.44	-0.52	0.44	134.56
Employment, 2002	5.22	3.32	2.00	1.49	-0.55	0.40	134.24

<i>Post weighting</i>							
Variable (in logs)	Mean		Variance		Skewness		Stand. Diff.
	Treated	Control	Treated	Control	Treated	Control	
TFP growth, 2004	0.03	0.03	0.05	0.05	0.01	0.01	0.00
TFP growth, 2003	-0.02	-0.02	0.04	0.04	0.05	0.05	0.00
Tangible fixed assets, 2004	6.16	6.16	3.91	3.91	-0.17	-0.17	0.00
Tangible fixed assets, 2003	6.14	6.14	3.96	3.96	-0.19	-0.19	0.00
Tangible fixed assets, 2002	6.11	6.11	3.97	3.97	-0.22	-0.22	0.00
Operating revenue, 2004	7.31	7.31	4.02	4.02	0.04	0.04	0.00
Operating revenue, 2003	7.26	7.26	4.07	4.07	0.02	0.02	0.00
Operating revenue, 2002	7.23	7.23	4.10	4.10	0.00	0.00	0.00
Employment, 2004	5.21	5.21	1.95	1.95	-0.46	-0.46	0.00
Employment, 2003	5.23	5.23	1.96	1.96	-0.52	-0.52	0.00
Employment, 2002	5.22	5.22	2.00	2.00	-0.55	-0.55	0.00

As indicated in Table 4, pre-balancing standardized differences lie far beyond these values for most covariates, in most cases even beyond 100. Balancing characteristics also indicate that EU ETS firms are substantially more sizeable than regulated firms with respect to their assets, revenue and employment. However, both groups do follow a common trend in productivity growth pre-treatment. While the assumption of a common trend under no treatment cannot be tested empirically, this observation does lend it some initial, albeit intuitive, support. As shown under "Post weighting", applying weights introduces a perfect balance between the two groups for each of the pre-2005 covariates and all three respective

moments. Standardized differences are 0 and each of the moments between the two groups is identical.

We now present our results obtained from estimating our model presented in Equation 2. For this purpose, we estimate a stylized version of the neo-Schumpeterian model, which incorporates innovation and productivity catch-up as two potential sources of firm's productivity growth, while at the same time accounting for persistent productivity dispersion within industries. This dynamic model allows us to differentiate the potential effects of the EU ETS on total factor productivity (TFP) growth depending on the level of firms' technological advancement.

Table 5 reports our main results. The first column shows the estimates of a basic pooled OLS model (1) that includes the full set of fixed effects and firm controls but does not account for dynamic adjustment effects. Similar to previous findings in the literature, we cannot actually find a significant effect of the EU ETS on TFP. Specification (2) introduces the important sources of technological progress and yields a negative effect of the EU ETS that is barely significant at the 10% level.

Table 5: Entropy balancing DiD

	Within-group			FD-GMM
	(1)	(2)	(3)	(4)
ETS treatment	-0.02 (0.013)	-0.03* (0.014)	0.06*** (0.015)	0.04*** (0.015)
gap TFP t-1 X ETS treatment			-0.09*** (0.016)	-0.09*** (0.023)
gap TFP t-1		0.44*** (0.040)	0.50*** (0.040)	0.25*** (0.068)
dlnTFP_Frontier		0.32*** (0.051)	0.33*** (0.049)	0.29*** (0.075)
Control variables	Yes	Yes	Yes	Yes
Time	Yes	Yes	Yes	Yes
Country x industry	Yes	Yes	Yes	-
Time x country	Yes	Yes	Yes	-
Firm	Yes	Yes	Yes	Yes
Observations	453.779	453.779	453.779	396.955
AB-AR(1)				0.000
AB-AR(2)				0.203
Hansen J (p-value)				0.132

Note: Standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

In line with the micro-econometric literature on the determinants of productivity growth (e.g. Griffith et al. (2004), Griffith et al. (2009), Conway et al. (2006), and Nicoletti and Scarpetta (2003)), we find strong evidence for the sources of technological advancement in all

specifications. The coefficient of TFP growth at the frontier is positive which points towards a pass-through of technologies developed at the frontier. We also find support for a catch-up process in the form of a positive coefficient for the gap term indicating that firms operating behind the frontier grow faster. Both effects are highly significant at the 1% level. Consistent with previous findings, the size of the coefficients highlights the importance of these two drivers of productivity growth.

In specification (3), we introduce the lagged catch-up term interacted with the DiD estimator and allow the effect of the EU ETS to vary depending on the given firm's position within the productivity distribution, i.e. the distance to the technology frontier. Results for specification 3 demonstrate that this step is crucial to properly identify the policy's impact. Both the simple DiD-term as well as the non-linear interaction effect are highly significant at the 1% level. The strong heterogeneity in firm responses to the policy makes clear why we could barely identify the effect in specification (2). The simple DiD-term stands for an on average increase in TFP growth by 6 percent and indicates positive effects for advanced firms, whereas the impact of the EU ETS in relation to the distance to the technology frontier points to an on average deceleration in TFP growth by 9%. Importantly, this finding clearly supports our hypothesis that the effect of the EU ETS on TFP growth is highly heterogeneous and depends on a firm's level of technological progress.

However, the introduction of firm fixed effects may give rise to a Nickell-bias in the estimates for the coefficient of the lagged gap term, which incorporates the lagged dependent variable. To address the simultaneity issue, we employ the Arellano-Bond estimator (4). While there are some slight differences in magnitude with respect to a somewhat smaller coefficient for the simple DiD term, the results are qualitatively identical when compared to our baseline results. These differences can also be due to the fact that the GMM set up does not allow us to include the full set of fixed effects. Note that, as expected, the simple gap-term is substantially smaller in magnitude which is due to the presence of a Nickell-bias and can also be related to the smaller number of fixed effects. However, the results indicate that, with regards to the policy impact of the EU ETS, our main findings are confirmed.

The challenge in the application of this GMM estimator is that with growing T the number of instruments can become large relative to sample size, which may render some asymptotic properties invalid (Roodman 2009). However, we successfully manage to limit the instrument count. We assess our instruments by using two standard test measures. The Hansen test has a null hypothesis of the instruments as a group being exogenous. The

corresponding p-value is clearly above the 10% significance level which suggests that our instruments are indeed exogenous. At the same time, it is not close to 1, which would point towards problems with respect to over-identification. Test results also suggest that standard error estimates are consistent in the presence of any pattern of heteroskedasticity and autocorrelation. The Arellano – Bond test has a null hypothesis of no autocorrelation and is applied to the differenced residuals. As it is common in such a setup, the first part of the test detects an AR(1) process in first differences. Hence, we instead check for an AR(2) process in first differences which would mean autocorrelation in levels. Autocorrelation in levels would indicate that lags of the dependent variable are in fact endogenous and thus bad instruments. Importantly, the corresponding p-value is clearly above the 10% significance level which leads us to confirm the hypothesis of no autocorrelation. Note that in a GMM setup, we cannot include a high number of fixed effects. In sum, our test results suggest that our findings with respect to the non-linear impact of the EU ETS are fully confirmed and not subject to a Nickell-bias.

Figure 1: Non-linear EU ETS impact on TFP growth

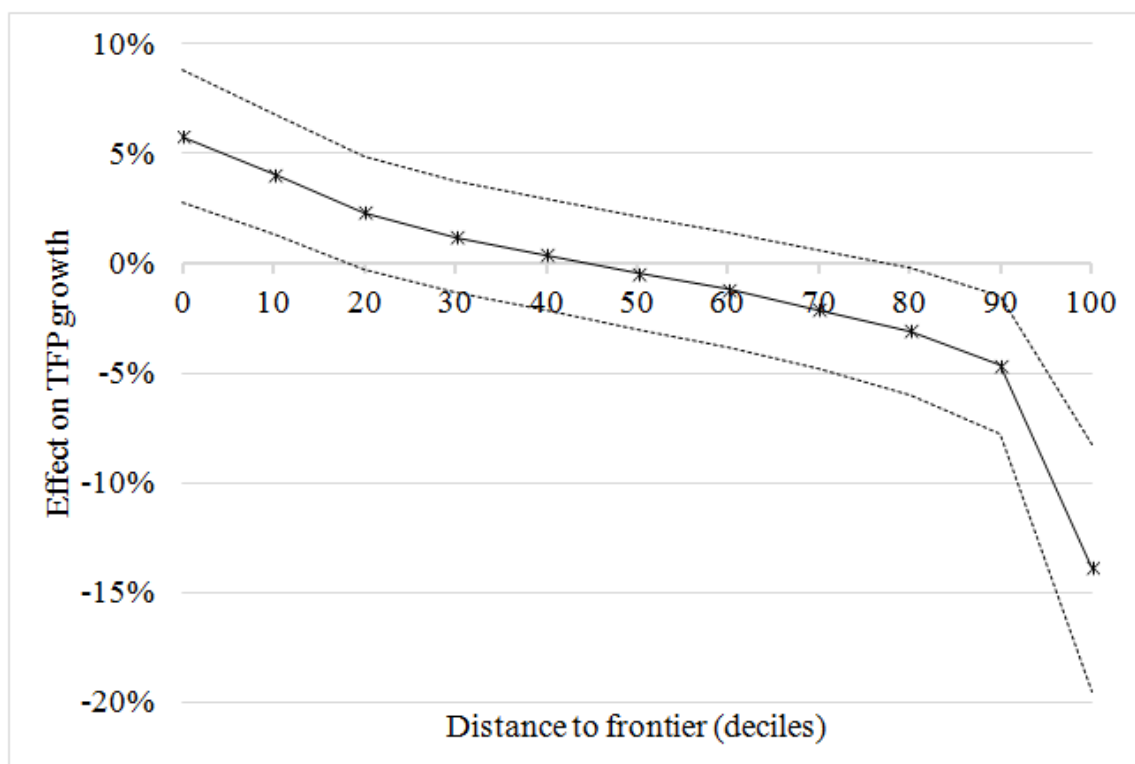


Figure 1 illustrates the estimated overall effect of the EU ETS on TFP growth with respect to a firm’s within-industry position. Based on Specification (3), the figure shows the

size of the joint EU ETS impact $(\beta_1 + \beta_2 \ln(\frac{A_{Fs}}{A_{ics}})_{t-1})$ in relation to the gap to the global frontier as shown in deciles of the respective distribution. We also plot the corresponding 95% confidence interval. The heterogeneity of the policy's effect seems to be apparent: The impact is positive for firms operating close to the frontier, subsequently declines in magnitude and significance, becomes insignificant for a sizeable group of firms, and eventually turns significant again but negative in terms of magnitude.

Whereas firms that operate very close to the frontier see their TFP growth accelerate by 5.8 percent, for firms at the 20th percentile the policy impact still corresponds to a 2.3 percent increase. The effect is highly significant at the 1% level until the 10th percentile and still significant at the 10% level for the 20th percentile. While a majority of firms appears to be unresponsive to the policy, firms operating beyond the 80th percentile start to see their productivity growth declining. At the 80th percentile, firms face a deceleration in TFP growth by 3.1 percent. The effect becomes gradually more negative in magnitude and drops sharply for firms operating far behind the frontier, signifying a 13.9 percent decrease at the 99th percentile. The negative section of the slope is highly significant (80th percentile: at the 5% level, 90th percentile and subsequently: at the 1% level).

In sum, the heterogeneity of the effect particularly suggests that technologically advanced firms benefit from the regulatory policy, whereas less-advanced EU ETS firms that operate further behind the productivity frontier are falling behind.

4.2 Robustness

We now employ a battery of tests to further assess the robustness of our findings and in particular with respect to the plausibility of our identification assumptions. The subsequent tests are based on or compared with Equation 2, i.e. the baseline specification with the full set of fixed effects (Table 5, Specification 3).

4.2.1 Omitted variable bias

If important unobserved covariates exist that both explain EU ETS participation and productivity, not including them will confound our estimate of the treatment effect of the EU ETS. For instance, if matched EU ETS firms had, even after our thorough balancing process, systematically higher innovative capacities than the control group, this could lead us to falsely attribute productivity growth linked with these characteristics to the regulatory policy.

We test the possibility of omitted variable bias and study coefficient stability as proposed by Oster (2019). Intuitively, the econometric estimation method compares changes in estimated coefficients and changes in the R^2 values of two different specifications with each other. The first contains a few basic explanatory variables only, and the second specification controls for a rich set of pre-determined explanatory variables. In case the R^2 in the full model is considerably higher than in the baseline model, this suggests that the additional explanatory variables help explaining variation in the outcome variable. If, at the same time, the estimated coefficient of interest does not change much from one model to the next, this suggests that the treatment effect is unlikely to be driven by unobserved characteristics. For a more detailed technical discussion we refer to Oster (2019).

One approach to implement this method is to calculate the so-called degree of proportionality δ , i.e. the explanatory power unobserved confounders would need relative to the explanatory power of observables to produce $\beta = 0$. For this purpose, one needs to make an assumption about the size of the R^2 of a (hypothetical) estimation that includes the full set of important observed and unobserved covariates (so-called R_{max}). Oster (2019) reviews the results from 27 articles in top journals that use either randomized or non-randomized data to obtain treatment effects and identifies $R_{max} = 1, 3R$ as an ideal bounding value. For this value, 90% of studies with randomized data and 45% of studies with nonrandomized data survive this test. Going beyond this bound or even up to $R_{max} = 1$ is not recommended as only 40% of studies with randomized data survive the test since even minimal changes in coefficients can be exaggerated under this assumption. Using $R_{max} = 1, 3R$, which implies that unobservables explain to some degree less than observables and treatment, also makes sense in our specific context as our observed covariates are important potential confounders and likely associated with other potential confounding factors.

Our test results are presented in Table 6 and imply that unobservables would need to be 2.9 times more important than other controls to explain away the non-linear treatment effect. The negative sign of the coefficient indicates that, if unobserved confounders indeed explained TFP to such a substantial degree, we would underestimate the negative impact and overestimate the positive impact of the EU ETS. However, this appears unlikely given our comprehensive set of control variables. In line with this argument, obtaining a value above 1 means that coefficient stability can be considered to be consistent with randomization (Oster 2019). In terms of the simple, positive treatment term, unobservable characteristics would need to be 2.4 times more important than our controls to explain away the

Table 6: Sensivity to omitted variable bias

	d for b = 0 given Rmax
ETS treatment	2.43
gap TFP t-1 X ETS treatment	2.91
R ²	0.2915
Rmax =1,3XR ²	0.3789
Control variables	Yes
Time	Yes
Country x industry	Yes
Time x country	Yes
Firm	Yes
Observations	453.779

Note: Based on Oster (2017) and baseline specification.

effect. In sum, these tests confirm that our results appear to be very robust with respect to potentially omitted confounders.

4.2.2 Matched DiD and propensity score weighted DiD

The results we obtain may be determined to some degree by the choice of technique we use to achieve covariate balance. To address this potential concern, we apply two alternative procedures. First, we employ a matched difference-in-differences ("Matched DiD") approach that utilizes propensity score matching. In contrast to the entropy balancing approach, propensity scores compress the information of the continuous pre-treatment variables into a single score (Rosenbaum and Rubin 1985). In practice, we obtain propensity scores for each firm in our sample by estimating a probit model that regresses regulatory status on up to three moments of the full set of potential confounders. The score is then used to match ETS firms with their respective closest neighbor from the reservoir of non-treated firms. Drawing from the overlap in propensity score distributions (so-called "common support") between ETS and non-ETS groups, we constrain the sample to those firms with a sufficiently close neighbor thus improving balance in covariate distributions. Treated firms for which we do not find a similar counterpart among non-EU ETS firms are discarded. This is a contrast to our entropy weighting approach that maintains a large sample size all while providing covariate balance. Hence, the main challenge when utilizing Matched DiD consists in designing a specification that balances out confounding factors without sacrificing too much sample size.

Second, we instead utilize the propensity scores to weight each firm-year-observation with its respective inverted probability to be treated, i.e. of being subject to the EU ETS

(Stuart 2010; Stuart et al. 2014). Applying inverted probability weights in this manner improves covariate balances between the treated and non-treated group while retaining a bigger sample size compared to a matching approach. We combine each of the two alternative balancing approaches with a difference-in-differences estimator to account for any remaining time-invariant differences. The alternative approaches also require firms to be matched within the same country-2 digit-industry pair. Enforcing a perfect balance in this regard accounts for potential unobserved differences such as sector-specific technologies while being more susceptible to measuring potential spillover effects as matched firm pairs are more likely to be competitors. Our approach of combining matching with DiD ("PS-Weighted DiD") follows Heckman et al. (1997) as summarized in Blundell and Dias (2009), whereas the approach of combining inverse-probability weighting with DiD follows Stuart (2010).

Table 7: Covariate balance after balancing, PS Matching and PS Weighting

Variable (in logs)	Stand. Diff.	
	PS-Matching	PS-Weighting
TFP growth, 2004	4.20	3.10
TFP growth, 2003	3.10	-0.70
Tangible fixed assets, 2004	-2.90	11.10
Tangible fixed assets, 2003	-2.60	11.70
Tangible fixed assets, 2002	-2.30	12.30
Operating revenue, 2004	-3.00	8.70
Operating revenue, 2003	-2.80	9.00
Operating revenue, 2002	-3.10	9.70
Employment, 2004	0.80	14.00
Employment, 2003	3.20	14.10
Employment, 2002	2.10	14.10
Exact matching, NACE 2 digit	Yes	Yes

Table 7 reports balancing results for the two procedures. In both cases, standardized differences indicate overall a good balance. For propensity score matching, the measure is well below ten for all covariates, whereas for the weighting procedure it is still below 15. Hence, covariate distributions between the two groups are now similar.

Table 8 presents the estimation results for the matched difference-in-differences estimator and the weighted difference-in-differences estimator along with the original baseline results. Compared to the baseline results, the findings are almost identical both in terms of magnitude and significance of the treatment and the non-linear term. This indicates that the EU ETS effects obtained from our original specification are robust to the choice of the balancing method. The results also demonstrate the major advantage of entropy weighting over the propensity score approach. While "Matched DiD" provides very good balance,

the sample size is clearly diminished. "PS-weighted DiD" maintains the sample size but at the expense of higher standardized differences. In contrast, entropy weighting keeps the original sample size (Table 5) all while establishing perfect balance (Table 4).

Table 8: Alternative Balancing Procedures

	Entropy balancing DID	PS Matched DID	PS Weighted DID
	(1)	(2)	(3)
ETS treatment	0.06*** (0.015)	0.05*** (0.017)	0.06*** (0.015)
gap TFP t-1 X ETS treatment	-0.09*** (0.016)	-0.09*** (0.016)	-0.10*** (0.015)
gap TFP t-1	0.50*** (0.040)	0.50*** (0.035)	0.53*** (0.035)
dlnTFP_Frontier	0.33*** (0.049)	0.33*** (0.049)	0.34*** (0.049)
Control variables	Yes	Yes	Yes
Time	Yes	Yes	Yes
Country x industry	Yes	Yes	Yes
Time x country	Yes	Yes	Yes
Firm	Yes	Yes	Yes
Observations	453.779	18.316	453.779
Exact match at NACE 2 digit level	-	Yes	Yes

Note: Standard errors in parentheses.* p<0.10, ** p<0.05, *** p<0.01.

4.2.3 Firm survival

Our results may also be subject to a bias arising from non-random differences in firm survival and exit mechanics between the two groups that are not related with the EU ETS (Griffith et al. 2009). Firm exits are often related with a previous slowdown in productivity growth as the given firm's capacity to compete declines (Albrizio et al. 2017). If non-EU ETS firms were more likely to exit the sample either by terminating their operations or by not reporting financial information for subsequent years, this would lead us to overestimate the positive and underestimate the negative treatment effect of the EU ETS. Vice versa, if EU ETS firms were less likely to survive, this would lead us to overestimate (underestimate) the negative (positive) impact of the EU ETS. Column 2 in Table 9 therefore presents estimation results for a sample of firms that exist throughout the entire period of analysis, i.e. report financial information in both 2002 and 2012. Both magnitudes and significance levels are very similar to our baseline results. This indicates that our results do not suffer from an attrition bias.

Table 9: Plausibility of unconfoundedness and stability of treatment value (SUTVA)

	Unconfoundedness			SUTVA
	Entropy balancing DID	Firm survival	Placebo effects	Late joiner
	(1)	(2)	(3)	(4)
ETS treatment	0.06*** (0.015)	0.06*** (0.016)		0.17 (0.103)
gap TFP t-1 X ETS treatment	-0.09*** (0.016)	-0.09*** (0.015)	-0.06 (0.052)	-0.10*** (0.031)
gap TFP t-1	0.50*** (0.040)	0.49*** (0.041)	1.22*** (0.034)	0.49*** (0.077)
dlnTFP_Frontier	0.33*** (0.049)	0.33*** (0.047)	0.39*** (0.039)	0.46*** (0.063)
Control variables	Yes	Yes	Yes	Yes
Time	Yes	Yes	Yes	Yes
Country x industry	Yes	Yes	Yes	Yes
Time x country	Yes	Yes	Yes	Yes
Firm	Yes	Yes	Yes	Yes
Observations	453.779	354.413	154.712	22.087

Note: Standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

4.2.4 Placebo effects

Another possibility to assess the plausibility of the unconfoundedness assumption is to apply a placebo analysis (Imbens and Rubin 2015). If our treatment effect was an artifact of systematically different omitted characteristics or underlying trends between the ETS and the non-ETS groups, one would expect us to detect a treatment effect even in the absence of actual treatment. It would also indicate the possibility of an announcement effect. The EU ETS directive was adopted at the end of 2003 before entering into force in 2005, a time span which could have given firms the possibility to adjust pre-policy. If an announcement effect had taken place, this may lead us to either over- or underestimate the impact of the EU ETS depending on if firms rather had incentives to select themselves into or out of the regulation. A placebo test thus allows us to address both of these concerns directly.

For this purpose, we restrict our sample to the years 2002-2004. We restructure our balancing procedure using exclusively covariates from the year 2002 and define for our difference-in-differences term the year 2003 as the start of the regulation instead of 2005. We redefine the regulatory dummy for ETS participation accordingly. Hence, operating revenue for the years 2003 and 2004 becomes our alternative pseudo-outcome. A priori we know that the pseudo-outcome should not have been affected as it was determined before the start of the policy. The corresponding effect for the main non-linear term as presented in Column 3 of Table 9 is unequivocally insignificant.⁷ This result lends further credibility to

⁷Note that given the short time period, we cannot include the non-linear effect jointly with the simple treatment term as the latter is omitted.

the assumption that important unobserved confounding factors are absent. In addition, it suggests that firms did not adjust in anticipation of the EU ETS as we detect no statistically different pattern in productivity dynamics between the two groups pre-treatment.

4.2.5 Stability of unit treatment value

We also test if our assumption of stability of unit treatment values (SUTVA), i.e. the absence of spillover effects between treatment and control group, is plausible under the given policy context. Given that our study is delivered at the firm level, we can rule out any potential within-firm spillovers that may occur between regulated installations (Petrick and Wagner 2014). However, non-EU ETS firms may still respond to the policy if, for instance, they compete with regulated firms. Untreated firms could then see their productivity affected by their competitors' policy-induced gains or losses. The SUTVA in itself cannot be assessed empirically. However, we can test specific institutional conditions under which a violation of the assumption would take place (Fowlie et al. 2012). In particular, if non-regulated firms indeed reacted to the regulation then we would expect a group of non-ETS firms which have been less exposed to the policy to respond less strongly (Calel and Dechezleprêtre 2016). In practice, we restrict the pool of control firms to firms from Norway, a country that did not join the EU ETS in 2005 but in 2008. Due to the small number of control firms, the test can only be conducted for a sample that requires less variables for the balancing process.

Column 4 of Table 9 shows that estimation results remain qualitatively similar when compared to our baseline results. Importantly, the main non-linear effect of the EU ETS is highly significant at the 1% level and very similar in magnitude. The coefficient for the simple treatment effect also retains its direction but is clearly more sizeable and barely insignificant (p-value: 0.110). However, this could be due to the fact that it is more challenging to identify the effect for the top percentiles of the productivity distribution in such a small sample. In sum, the result for the "late joiner test" does not lend support to the idea that our estimate of the ATET might be an artifact of control firms reacting to the policy.

The impact of the EU ETS is highly significant and non-linear in nature, i.e. the effect becomes more negative the further a firm operates from the productivity frontier. If anything, the results suggest that the effect for firms operating close to the frontier might be less negative but not necessarily positive. However, due to between country differences for which these estimates cannot control, we have to be cautious with this interpretation. The

re-balanced estimates do not differ significantly from our original estimates. Hence, they do not represent a substantive challenge to our findings.

4.3 Heterogeneous treatment effects

We now assess the potential channels of the non-linear effect of the EU ETS in-depth by relying on the rich information on firm-, country-, and industry-specific characteristics of our dataset. Our assumption is that magnitude and significance of the two-directional effect may be more pronounced for certain subgroups of firms. We categorize the sample into two sub-samples, one of Northern and one of Southern European countries, and re-estimate entropy balancing weights for both of these samples.⁸ Subsequent results are qualitatively identical when we only restrict ETS firms to a certain region while retaining firms from all 8 European countries as the pool of respective control firms.

Table 10: Heterogeneity

	Entropy balancing DID	Northern Europe	Southern Europe
	(1)	(2)	(3)
ETS treatment	0.06*** (0.015)	0.09*** (0.021)	0.06*** (0.015)
gap TFP t-1 X ETS treatment	-0.09*** (0.016)	-0.10*** (0.023)	-0.09*** (0.016)
gap TFP t-1	0.50*** (0.040)	0.44*** (0.058)	0.50*** (0.040)
dlnTFP_Frontier	0.33*** (0.049)	0.32*** (0.037)	0.33*** (0.049)
Control variables	Yes	Yes	Yes
Time	Yes	Yes	Yes
Country x industry	Yes	Yes	Yes
Time x country	Yes	Yes	Yes
Firm	Yes	Yes	Yes
Observations	453.779	73.708	453.779


Note: Standard errors in parentheses.* p<0.10, ** p<0.05, *** p<0.01.

As shown in Table 10, estimation results for both subsamples confirm a non-linear effect. However, with respect to the Northern sample of EU countries (Figure 2), the positive section of the non-linear effect is considerably more pronounced both in terms of magnitude and significance when compared to our baseline results. In addition, a wider range of firms appears to benefit from the EU ETS. Whereas firms operating very close to the frontier see their TFP growth accelerate by 8.9 percent, even for firms at the 40th percentile the effect still corresponds to a 3.8 percent increase. The effect is highly significant at the 1% level until the

⁸The countries contained in the “Northern sample” are Belgium, France, Germany, Great Britain, Norway and Sweden. The countries in the “Southern sample” are Italy and Spain.

20th and significant at the 5% level up until the 40th percentile. In contrast, the negative effect becomes significant only beyond the 90th percentile, whereas the large remainder of firms is unresponsive. For such a firm, TFP growth can still decelerate by up to 16.5 percent.

Figure 2: Northern countries: Non-linear EU ETS impact on TFP growth




Results for the sample of Southern European countries (Figure 3) draw a different picture. Although the shape of the magnitude remains non-linear, the vast majority of firms is unresponsive and firms operating at and beyond the 70th percentile face a deceleration in TFP growth that ranges between 2.6 and up to 7.8 percent. After and including the 90th percentile, the effect becomes highly significant at the 1% level (70th percentile: 10%, 80th percentile: 5%).

5 Discussion

Empirical evidence on the economic impacts of the EU Emissions Trading System is important to inform policy-makers but it has largely remained inconclusive. This study contributes to this discussion by investigating the potential effect of the EU ETS on total factor productivity growth as a summary estimate of the costs and benefits borne by firms.

Figure 3: Southern countries: Non-linear EU ETS impact on TFP growth



Our results suggest that the effect of the EU ETS is highly heterogeneous and depends on the distance to the technological frontier. This brings together two central paradigms on the economic impacts of environmental regulation: Negative effects for technologically less advanced firms speak to the conventional wisdom that policy impacts can entail efficiency decreases as a result of factor reallocation. In contrast, the finding that technologically more advanced firms increase their productivity growth lends support to the Porter Hypothesis that environmental policy induces efficiency gains as firms benefit from innovations or adoption of more efficient technology. We demonstrate that results are robust to the potential omission of confounding factors, sample attrition and different choices in estimators and covariate balancing techniques. Test results also point towards the absence of announcement effects and spillovers between treated and untreated firms.

This has important implications that can be taken into consideration for future policy design. Firms can substantially benefit from market based climate policy and increase their competitiveness. Altering the productivity distribution also entails distributional consequences which, in a Neo-Schumpeterian sense, may be beneficial if inefficient firms are pushed out of the market. At the same time, if the productivity gap consistently widens for less advanced firms, this could undermine their ability to compete in international mar-

kets. However, several studies demonstrate that the EU ETS has induced firms to substantially invest into their European asset bases without showing any signs of relocating to non-regulated world regions (see Moore et al. 2019; Koch and Basse-Mamba 2019). If negative effects for some firms indeed persisted, this could give rise to auxiliary policies aimed at firms transitioning more smoothly towards low-carbon technologies or enabling them to more rapidly adopt the benefits generated at the frontier. Hence, in terms of future research it would be interesting to see how benefits and costs develop in future regulatory phases, in particular under higher allowance prices. Heterogeneity in firms' ability to respond to carbon regulation may be particularly important in the context of a cross-country policy.

Our findings also indicate that environmental policy can exacerbate pre-existing structural differences in productivity dynamics among European economies. While benefits in terms of productivity growth are very pronounced in Northern European countries and broadly shared among a large number of firms, a higher degree of firms is either unresponsive or more prone to negative effects in Southern European countries. This observation is consistent with previous studies on persistent heterogeneity in the European productivity distribution (Verschelde et al. 2016). This could underline the case for policies to improve country-specific institutional conditions under which these firms are operating.

References

- Abrell, J., A. Ndoye, and G. Zachmann (2011). *Assessing the impact of the EU ETS using firm level data*. Bruegel Working Paper 2011/08.
- Akerberg, D. A., K. Caves, and G. Frazer (2015). Identification Properties of Recent Production Function Estimators. *Econometrica* 83 (6), 2411–2451.
- Aghion, P. and P. Howitt (2006). Appropriate Growth Policy: A Unifying Framework. *Journal of the European Economic Association* 4 (2-3), 269–314.
- Albrizio, S., T. Kozluk, and V. Zipperer (2017). Environmental policies and productivity growth: Evidence across industries and firms. *Journal of Environmental Economics and Management* 81, 209–226.
- Ambec, S., M. A. Cohen, S. Elgie, and P. Lanoie (2013). The Porter Hypothesis at 20: Can Environmental Regulation Enhance Innovation and Competitiveness? *Review of Environmental Economics and Policy* 7 (1), 2–22.
- Andrews, D., C. Criscuolo, and P. N. Gal (2015). Frontier Firms, Technology Diffusion and Public Policy: Micro Evidence from OECD Countries.
- Arellano, M. and S. Bond (1991). Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations. *The Review of Economic Studies* 58 (2), 277–297.
- aus dem Moore, N., P. Großkurth, and M. Themann (2019). Multinational corporations and the EU Emissions Trading System: The specter of asset erosion and creeping deindustrialization. *Journal of Environmental Economics and Management* 94, 1–26.
- Bartelsman, E. J. and M. Doms (2000). Understanding Productivity: Lessons from Longitudinal Microdata. *Journal of Economic Literature* 38 (3), 569–594.
- Becker, R. A. (2011). Local environmental regulation and plant-level productivity. *Ecological Economics* 70 (12), 2516–2522.
- Becker, R. A. et al. (2006). Micro and Macro Data Integration: The Case of Capital. In: *NBER Chapters*. National Bureau of Economic Research, Inc, 541–610.
- Bertrand, M., E. Duflo, and S. Mullainathan (2004). How Much Should We Trust Differences-In-Differences Estimates? *The Quarterly Journal of Economics* 119 (1), 249–275.

- Blundell, R. and M. C. Dias (2009). Alternative approaches to evaluation in empirical microeconomics. *Journal of Human Resources* 44 (3), 565–640.
- Bourlès, R. et al. (2012). Do Product Market Regulations In Upstream Sectors Curb Productivity Growth? Panel Data Evidence For Oecd Countries. *The Review of Economics and Statistics* 95 (5), 1750–1768.
- Calel, R. and A. Dechezleprêtre (2016). Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market. *Review of Economics and Statistics* 98 (1), 173–191.
- Chan, H. S., S. Li, and F. Zhang (2013). Firm competitiveness and the European Union emissions trading scheme. *Energy Policy* 63 (C), 1056–1064.
- Cohen, M. A. and A. Tubb (2017). The Impact of Environmental Regulation on Firm and Country Competitiveness: A Meta-analysis of the Porter Hypothesis. *Journal of the Association of Environmental and Resource Economists* 5 (2), 371–399.
- Cole, M. A. and R. J. R. Elliott (2003). Do Environmental Regulations Influence Trade Patterns? Testing Old and New Trade Theories. *The World Economy* 26 (8), 1163–1186.
- Collard-Wexler, A. and J. De Loecker (2016). *Production Function Estimation with Measurement Error in Inputs*. Working Paper 22437. National Bureau of Economic Research.
- Commins, N., S. Lyons, M. Schiffbauer, and R. S. Tol (2011). Climate Policy & Corporate Behavior. *The Energy Journal* 32 (4), 51–68.
- Conway, P., D. De Rosa, G. Nicoletti, and F. Steiner (2006). *Regulation, Competition and Productivity Convergence*. OECD Economics Department Working Paper 509. OECD Publishing.
- Dechezleprêtre, A. and M. Sato (2017). The Impacts of Environmental Regulations on Competitiveness. *Rev Environ Econ Policy* 11 (2), 183–206.
- Fowlie, M., S. P. Holland, and E. T. Mansur (2012). What do emissions markets deliver and to whom? Evidence from Southern California’s NOx trading program. *American Economic Review* 102 (2), 965–993.
- Gal, P. N. (2013). Measuring Total Factor Productivity at the Firm Level using OECD-ORBIS.
- Garrido, M. M. et al. (2014). Methods for Constructing and Assessing Propensity Scores. *Health Services Research* 49 (5), 1701–1720.

- Gopinath, G., Kalemli-Özcan, L. Karabarbounis, and C. Villegas-Sanchez (2017). Capital Allocation and Productivity in South Europe. *Q J Econ* 132 (4), 1915–1967.
- Gray, W. B. (1987). The Cost of Regulation: OSHA, EPA and the Productivity Slowdown. *The American Economic Review* 77 (5), 998–1006.
- Gray, W. B. and R. J. Shadbegian (2001). *Plant Vintage, Technology, and Environmental Regulation*. Working Paper 8480. National Bureau of Economic Research.
- Greenstone, M., J. List, and C. Syverson (2012). *The Effects of Environmental Regulation on the Competitiveness of U.S. Manufacturing*. Cambridge and MA: National Bureau of Economic Research.
- Griffith, R., S. Redding, and J. V. Reenen (2004). Mapping the Two Faces of R&D: Productivity Growth in a Panel of OECD Industries. *The Review of Economics and Statistics* 86 (4), 883–895.
- Griffith, R., S. Redding, and H. Simpson (2009). Technological Catch-up and Geographic Proximity*. *Journal of Regional Science* 49 (4), 689–720.
- Hainmueller, J. (2012). Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies. *Polit. anal.* 20 (1), 25–46.
- Hainmueller, J. and Y. Xu (2013). **ebalance** : A Stata Package for Entropy Balancing. *J. Stat. Soft.* 54 (7).
- Harris, M. N., L. Kónya, and L. Mátyás (2002). Modelling the Impact of Environmental Regulations on Bilateral Trade Flows: OECD, 1990–1996. *The World Economy* 25 (3), 387–405.
- Heckman, J. J., H. Ichimura, and P. E. Todd (1997). Matching As An Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme. *The Review of Economic Studies* 64 (4), 605–654.
- Imbens, G. W. and D. B. Rubin (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. 1st ed. Cambridge University Press.
- Jaffe, A. B. and K. Palmer (1997). Environmental Regulation and Innovation: A Panel Data Study. *The Review of Economics and Statistics* 79 (4), 610–619.
- Jaffe, A. B., S. R. Peterson, P. R. Portney, and R. N. Stavins (1995). Environmental Regulation and the Competitiveness of U.S. Manufacturing: What Does the Evidence Tell Us? *Journal of Economic Literature* 33 (1), 132–163.

- Kalemli-Ozcan, S. et al. (2015). *How to Construct Nationally Representative Firm Level data from the ORBIS Global Database*. Working Paper 21558. National Bureau of Economic Research.
- Koch, N. and H. Basse Mama (2019). Does the EU Emissions Trading System induce investment leakage? Evidence from German multinational firms. *Energy Economics* 81, 479–492.
- Kozluk, T. and V. Zipperer (2015). Environmental policies and productivity growth: a critical review of empirical findings. *OECD Journal: Economic Studies* 2014 (1), 155–185.
- Lechner, M. (2010). The Estimation of Causal Effects by Difference-in-Difference Methods Estimation of Spatial Panels. *Foundations and Trends® in Econometrics* 4 (3), 165–224.
- Levinsohn, J. and A. Petrin (2003). Estimating Production Functions Using Inputs to Control for Unobservables. *The Review of Economic Studies* 70 (2), 317–341.
- Lutz, B. J. (2016). *Emissions trading and productivity: Firm-level evidence from German manufacturing*. Tech. rep. 16-067. ZEW - Leibniz Centre for European Economic Research.
- Löschel, A., B. J. Lutz, and S. Managi (2019). The impacts of the EU ETS on efficiency and economic performance – An empirical analyses for German manufacturing firms. *Resource and Energy Economics*. Recent Advances in the Economic Analysis of Energy Demand - Insights for Industries and Households 56, 71–95.
- Marin, G., M. Marino, and C. Pellegrin (2018). The Impact of the European Emission Trading Scheme on Multiple Measures of Economic Performance. *Environmental & Resource Economics* 71 (2), 551–582.
- McConnell, V. D. and R. M. Schwab (1990). The Impact of Environmental Regulation on Industry Location Decisions: The Motor Vehicle Industry. *Land Economics* 66 (1), 67–81.
- Morgenstern, R. D., W. A. Pizer, and J.-S. Shih (2002). Jobs Versus the Environment: An Industry-Level Perspective. *Journal of Environmental Economics and Management* 43 (3), 412–436.
- Nickell, S. (1981). Biases in Dynamic Models with Fixed Effects. *Econometrica* 49 (6), 1417.
- Nicoletti, G. and S. Scarpetta (2003). Regulation, productivity and growth: OECD evidence. *Econ Policy* 18 (36), 9–72.
- Olley, G. S. and A. Pakes (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica* 64 (6), 1263–1297.

- Oster, E. (2019). Unobservable Selection and Coefficient Stability: Theory and Evidence. *Journal of Business & Economic Statistics* 37 (2), 187–204.
- Petrick, S. and U. J. Wagner (2014). *The Impact of Carbon Trading on Industry: Evidence from German Manufacturing Firms*. Kiel Working Papers 1912. Kiel Institute for the World Economy.
- Porter, M. E. and van der Linde, Claas (1995). Toward a New Conception of the Environment-Competitiveness Relationship. *The Journal of Economic Perspectives* 9 (4), 97–118.
- Roodman, D. (2009). A Note on the Theme of Too Many Instruments*. *Oxford Bulletin of Economics and Statistics* 71 (1), 135–158.
- Rosenbaum, P. R. and D. B. Rubin (1985). Constructing a Control Group Using Multivariate Matched Sampling Methods That Incorporate the Propensity Score. *The American Statistician* 39 (1), 33.
- Stuart, E. A. (2010). Matching Methods for Causal Inference: A Review and a Look Forward. *Statist. Sci.* 25 (1), 1–21.
- Stuart, E. A. et al. (2014). Using propensity scores in difference-in-differences models to estimate the effects of a policy change. *Health Serv Outcomes Res Method* 14 (4), 166–182.
- Syverson, C. (2011). What Determines Productivity? *Journal of Economic Literature* 49 (2), 326–365.
- Vershelde, M., M. Dumont, G. Rayp, and B. Merlevede (2016). Semiparametric stochastic metafrontier efficiency of European manufacturing firms. *Journal of Productivity Analysis* 45 (1), 53–69.
- Wooldridge, J. M. (2009). On estimating firm-level production functions using proxy variables to control for unobservables. *Economics Letters* 104 (3), 112–114.