



HAL
open science

Estimating the Effects of Regulation When Treated and Control Firms Compete: A New Method with Application to the EU ETS

Geoffrey Barrows, Raphael Calel, Martin Jégard, Hélène Ollivier

► **To cite this version:**

Geoffrey Barrows, Raphael Calel, Martin Jégard, Hélène Ollivier. Estimating the Effects of Regulation When Treated and Control Firms Compete: A New Method with Application to the EU ETS. 2024. <hal-04378537>

HAL Id: hal-04378537

<https://hal.science/hal-04378537v1>

Preprint submitted on 8 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

**Estimating the Effects of
Regulation When Treated and
Control Firms Compete:
A New Method with
Application to the EU ETS**

Geoffrey Barrows, Raphael Calel, Martin Jégard, Hélène Ollivier

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Estimating the Effects of Regulation When Treated and Control Firms Compete: A New Method with Application to the EU ETS

Abstract

This paper presents a method for estimating treatment effects of regulations when treated and control firms compete on the output market. We develop a GMM estimator that recovers reduced-form parameters consistent with a model of differentiated product markets with multi-plant firms, and use these estimates to evaluate counterfactual revenues and emissions. Our procedure recovers unbiased estimates of treatment effects in Monte Carlo experiments, while difference-in-differences estimators and other popular methods do not. In an application, we find that the European carbon market reduced emissions at regulated plants without undermining revenues of regulated firms, relative to an unregulated counterfactual.

JEL-Codes: Q480, L100, L500.

Keywords: regulation, spillovers, environment, energy, firms.

Geoffrey Barrows
Institut Polytechnique de Paris
France – 91120 Palaiseau
geoffrey-masters.barrows@polytechnique.edu

Raphael Calel
McCourt School of Public Policy, Georgetown
University / Washington DC / USA
raphael.calel@georgetown.edu

Martin Jégard
Swedish University of Agricultural Sciences
Economics Department
Uppsala / Sweden
martin.jegard@slu.se

Hélène Ollivier
Paris School of Economics – CNRS
48 Boulevard Jourdan
France – 74014 Paris
helene.ollivier@psemail.eu

We thank Antoine Dechezleprêtre, Robert Elliot, Elia Lapenta, Isabelle Méjean, Erich Muehlegger, Ariell Reshef, Gregor Singer, Farid Toubal, Francesco Vona, and seminar participants at École Polytechnique, CEPII, The Frankfurt School of Finance and Management, CESifo, Sustainable Development, Environment and Energy seminar of the University of Paris-Nanterre, Power to Empower Emerging Africa, The Toulouse School of Economics, EEME workshop Zurich, PSE Annual Conference on Global Issues, and The Occasional Workshop in Environmental and Resource Economics UCSB for useful comments and discussion. Barrows acknowledges support from the grant Investissements d’Avenir (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047). Ollivier acknowledges support from the EUR grant ANR-17-EURE-0001.

1 Introduction

Governments often regulate particular firms differently from other firms in the same industry. For example, governments impose stricter environmental standards for firms in regions with higher pollution (Greenstone, 2002; Fowlie et al., 2016; Martin et al., 2014), incentivize employment for firms in distressed neighborhoods (Neumark & Simpson, 2015), and offer special protections and subsidies for smaller firms (Martin et al., 2017; Rotemberg, 2019). When evaluating the effects of these policies, researchers frequently compare the evolution of outcomes for regulated firms to those of unregulated firms within the same industry, controlling flexibly for common trends in input prices, productivity, and demand. But if firms compete in the output market, then virtually any model of imperfect competition would predict that the effects of industrial regulation spill over from regulated to unregulated firms through output prices. This core insight from microeconomic theory implies that conventional difference-in-differences (DD) estimators fail to identify treatment effects, as the outcomes of control firms are contaminated by the very policy of interest.¹

In this paper, we propose a new method for estimating treatment effects of an incomplete regulation in the presence of interfirm spillovers. We start by specifying a model of supply and demand for differentiated goods in which the government taxes the inputs of a subset of firms in the economy. We rely on standard assumptions of producer and consumer behaviors from industrial organization and international trade theory – CES demand, Cobb-Douglas production, heterogeneous productivities, and monopolistic competition.² We show that, in this framework, the evolution of firm-level inputs and outputs between a pre-regulation period and a post-regulation period depends on a firm’s own regulation status (through unit costs), as well as on the regulation status of all firms in the sector (through the industry- and sector-wide price indices). Ignoring these across-firm dependencies leads to biased estimates of average and aggregate treatment effects. To overcome this problem, we develop a Generalized Method of Moments (GMM) procedure that leverages model-based exclusion restrictions to estimate reduced-form coefficients. We

¹Formally, across-firm dependencies represent a violation of the stable unit treatment value assumption (SUTVA), which is a necessary assumption in DD estimation, and more generally, two-way fixed effect models.

²These assumptions are standard for models of the manufacturing sector as a whole (see for example Atkeson & Burstein 2008 and Shapiro & Walker 2018). When researchers study a single industry, flexible patterns of across-firm price elasticities can be estimated structurally using detailed output price and product characteristic information (Berry et al., 1995). However, data of this sort are only available for specific industries, hence this method is not well-suited to estimating aggregate effects of regulation (De Loecker & Syverson, 2021).

show that estimates of these reduced-form coefficients are sufficient to compute counterfactual outcomes, and hence, average and aggregate treatment effects of the policy.³

We use our method to provide a new assessment of the European Union’s Emissions Trading System’s (EU ETS) effect on French manufacturers’ revenues and CO₂ emissions. The EU launched the world’s first major cap-and-trade program for carbon emissions in 2005, regulating the emissions of over 10,000 producers across all manufacturing sectors. A key feature of the EU ETS is that it only regulates large emitting installations, which means that regulation typically varies within industries, and even across plants within a firm. Given this feature, we explicitly model the behavior of multi-plant firms in order to derive a microfounded method for aggregating plant-level regulation assignments to the firm level. We then implement our new GMM procedure using administrative data on French manufacturers’ revenues and CO₂ emissions and compare to existing estimation methods.

Analytically, we show that the regulation’s true effects on revenues and emissions cannot, in general, be signed based on the underlying structural parameters alone. The effect will depend on how parameters, like the endogenous technological response of regulated firms, interact with the pre-existing distribution of market shares. Moreover, we show that the DD estimates can be biased towards zero or away from zero, and can even deliver the wrong sign, in expectation. The common intuition that the DD estimator exaggerates the effect of regulation on firm-level revenues holds only in a world exclusively populated by single-plant firms.

Given our data generating process, we can also evaluate the performance of existing alternatives to conventional DD estimators. A handful of papers augment the conventional DD estimation equation with controls for treatment density within a market, where a market is usually defined by a contiguous geographical unit (Cai & Szeidl, 2022; Muehlegger & Sweeney, 2021; Rotemberg, 2019). We show that this approach amounts to taking a local approximation around the pre-treatment equilibrium – henceforth, we refer to this as the local approximation (LA) estimator. This approximation is valid as long as market shares do not change much over time. We derive an analytical expression for the bias in

³In some contexts, when structural parameters of supply and demand are available in the literature, it may be possible to estimate treatment effects using computable general equilibrium models. However, in this case, the structural assumptions imposed in estimating parameters may not be consistent with the assumptions required for computing counterfactuals. Additionally, oftentimes parameters are estimated on samples that differ from the context of the study, and hence may not be valid. In the method we propose, all parameters can be estimated from the data in a way that is consistent with standard microfoundations that are made explicit.

the general case, and document in Monte Carlo experiments that this local approximation can lead to biases in either direction. By contrast, we document that our GMM procedure delivers unbiased estimates of treatment effects even in finite samples. Relative to the local approximation method, our approach also generalizes by allowing for non-random assignment of regulation, non-random-walk productivity growth, and imported goods.

Empirically, we find that, even though the regulation raised the cost of energy, the EU ETS did not disadvantage regulated firms. On the contrary, we estimate that regulated firms increased their annual sales between 6-9% as a result of the EU ETS. Critically, we find evidence of interfirm spillover effects of the regulation – the revenue growth of individual firms was affected by the density of regulated firms in the same industry. At the same time, we find that CO₂ emissions fell between 5-25% on average at regulated plants, depending on the year. These results are consistent with the Porter hypothesis: by investing in emissions-reducing technologies, regulated firms lowered emissions and lowered costs at the same time (Porter & van der Linde, 1995; Ambec et al., 2013). This mechanism is also supported by evidence that the EU ETS has spurred greater R&D and green innovation among regulated firms (Calel & Dechezlepretre, 2016; Calel, 2020).

In aggregate, we find that the EU ETS reduced CO₂ emissions generated by French manufacturers in the production of goods for the domestic market by 0.9-4.6 million tonnes annually, or between 3-16% of observed domestic emissions, relative to an unregulated counterfactual. By contrast, we find that the DD estimator overstates the effect on revenues and understates the effect on plant-level and aggregate emissions, while the LA estimator understates the effect on revenues and overstates the effect on emissions.

This paper contributes to at least four strands of research. First, we contribute to recent literature on the estimation of treatment effects with DD techniques. Several recent papers demonstrate that DD estimators – and more broadly, two-way fixed effect estimators – may yield biased estimates of average treatment effects when a policy’s effect is heterogeneous across groups or over time (see de Chaisemartin & D’Haultfoeuille 2022 for a review). Like this heterogeneous treatment effect literature, we present a new method for estimating treatment effects when an assumption of the DD framework fails – SUTVA, in our case.

Second, we contribute to a small group of papers that explicitly account for across-unit spillovers when estimating treatment effects (Borusyak et al., 2022; Adao et al., 2020; Franklin et al., 2021; Bergquist et al., 2022). These papers mostly study the effect of demand or technology shocks on local labor supply in the context of mobile workers. By focusing on output market competition, we study a different mechanism. Our paper is

similar in spirit to Borusyak et al. (2022), who document the bias in conventional DD estimates of the elasticity of migration to demand shocks, and then propose a model-based approach.

Third, we contribute to a literature that structurally estimates the effects of environmental policies on emissions and competitiveness (Fowlie et al., 2016; Muehlegger & Sweeney, 2021; Hintermann, 2017; Ganapati et al., 2020; Fabra & Reguant, 2014). Most of this literature studies homogeneous product markets and often focuses on pass-through from input prices to output prices. A notable exception is Shapiro & Walker (2018), who study the effects of environmental regulations on aggregate emissions from US manufacturing. Our model shares many features with theirs, but the context and empirical strategy are quite different. Shapiro & Walker (2018) back out the implied homogeneous environmental tax from aggregate data on sector-level emissions and revenues, while we study the effect of a particular (incomplete) regulation on a discrete set of firms. Since the regulation status of particular firms matters in our context, we solve counterfactuals firm by firm, while Shapiro & Walker (2018) aggregate over firms in the industry to solve for counterfactuals.

Finally, our paper contributes to the literature evaluating the effect of the EU ETS on firm-level emissions and revenues. Most of the studies so far have used DD estimators at the firm- or plant-level, comparing firms with a regulated plant to those without (Löschel et al., 2019; Colmer et al., 2021; Dechezleprêtre et al., 2018; Jaraite & Di Maria, 2016).⁴ Our analysis uncovers evidence of interfirm spillover effects of the regulation, however, which violates SUTVA.⁵ Our analytical results demonstrate that, under these conditions, the DD estimator cannot in general identify either the sign or the magnitude of the treatment effects. In the special case of single-plant firms, the DD estimator can only recover the sign of the effect on revenues. Our empirical findings, then, even where they appear similar to earlier estimates, provide a more credible basis for assessing the effects of the European carbon market.

⁴See Martin et al. (2015) and Joltreau & Sommerfeld (2018) for reviews of the literature.

⁵Other studies of the EU ETS investigate productivity and cost pass-through directly using production function estimation techniques (Calligaris et al., 2022). In most models, productivity is not conditioned by the choices of other firms in the economy, in which case the SUTVA holds. Still, revenues and emissions are determined in equilibrium, and hence depend on the strategic interaction of firms. While regulation may have important effects on productivity and pass-through, it is still necessary to model and estimate the demand side of the market in order to compute treatment effects on revenues and emissions.

2 Model

This section presents a model of supply and demand for differentiated goods in which the government regulates production at a subset of firms in the economy. Given our empirical application, we refer to the regulated output as “emissions,” but the model could be used to study the effects of a wide array of industrial policies and shocks. We initially describe an economy populated by single-plant firms and derive expressions for the effects of an incomplete regulation on revenues and emissions. We then extend the model to include multi-plant firms, and consider a regulation that targets a subset of plants in the economy.

2.1 Demand

Consider an economy in which a representative consumer divides expenditures between a set of differentiated products, which we refer to as “manufacturing” products, and a single homogeneous good, which we refer to as the “outside” good. The representative consumer’s preferences over varieties from manufacturing sectors $s = \{1, \dots, S\}$ and the outside sector $s = 0$ are described by the following three-tiered utility function:

$$U_t = (Q_{0t})^{a_0} \prod_s \left[\left(\sum_{i \in \Upsilon_s} (Q_{ist})^\nu \right)^{1/\nu} \right]^{a_s} \text{ where } Q_{ist} = \left(\sum_{f \in \Omega_{ist}} (Q_{fist})^\rho \right)^{1/\rho} \quad (1)$$

where Q_{0t} , Q_{ist} , and Q_{fist} denote the consumption at time t of the outside good, aggregate consumption in industry i , and consumption of individual variety f , respectively. The first tier aggregates consumption in a Cobb-Douglas function across sectors, which implies that expenditures on each sector s , Y_{st} , are determined as fixed shares of total expenditures, Y_t : $Y_{st} = a_s Y_t$. The second and third tiers aggregate consumption via Constant Elasticity of Substitution (CES) functions across the set of industries within a sector s , Υ_s , and across the set of varieties available in each industry i at time t , Ω_{ist} . We assume varieties are imperfect substitutes within an industry, industry-wide aggregates are imperfect substitutes within a sector, and varieties within an industry are closer substitutes than varieties from other industries, implying $0 < \nu < \rho < 1$.⁶

Utility maximization implies that expenditures on variety f at time t can be written

⁶This demand system (or very near versions of it) is a workhouse model of consumer demand in industrial organization and international trade theory (see for example Shapiro & Walker 2018 and Atkeson & Burstein 2008). An important feature of this demand system for our purpose is that it is general enough to allow revenue gains from one set of firms to crowd out sales of other firms in the same industry, though the crowding out is not necessarily complete.

as a function of the price of variety f at time t , p_{fist} , the price index for industry i , P_{ist} , the price index for sector s , Ψ_{st} , and sector-wide expenditures:

$$y_{fist} = (p_{fist})^{\frac{\rho}{\rho-1}} (P_{ist})^{\frac{\nu-\rho}{(\nu-1)(1-\rho)}} (\Psi_{st})^{\frac{\nu}{1-\nu}} Y_{st} \quad (2)$$

with

$$P_{ist} \equiv \left(\sum_{k \in \Omega_{ist}} (p_{kist})^{\frac{\rho}{\rho-1}} \right)^{\frac{\rho-1}{\rho}}, \quad \Psi_{st} \equiv \left(\sum_{m \in \Upsilon_s} (P_{mst})^{\frac{\nu}{\nu-1}} \right)^{\frac{\nu-1}{\nu}}. \quad (3)$$

Given the parameter restrictions, expenditures on variety f from industry i and sector s are decreasing in variety f 's own price, but increasing as a function of industry i 's price index, sector s 's price index, and sector-wide expenditures.

2.2 Production and Emissions

Each manufacturing firm produces a single differentiated variety, so f can be used interchangeably to index both varieties and firms. Manufacturing firms combine two variable inputs – labor, L , and energy, E – to produce output, Q , using Cobb-Douglas technology with constant returns to scale.⁷

Firms are heterogeneous with respect to productivity, which we allow to evolve over time according to a flexible Markov process:

$$\omega_{fist} = g(\omega_{fist-1}) + u_{fist}, \quad (4)$$

where $g(\cdot)$ is an arbitrary function of past firm-level productivity and u_{fist} is an i.i.d. shock with mean zero. If $g(\omega_{fist-1}) = \omega_{fist-1}$, then productivity follows a random walk. Otherwise, productivity is path dependent. If $g'(\omega_{fist-1}) < 1$, more productive firms grow slower, whereas more productive firms grow faster if $g'(\omega_{fist-1}) > 1$.

Production generates emissions, Z , in proportion to the amount of energy input: $Z_{fist} = \kappa_t E_{fist}$, where κ_t is the quantity of emissions per unit of energy, which may vary over time.

⁷The Cobb-Douglas assumption ensures that the effects of the regulation on input prices and the productivity combine linearly in the expression for revenues. If the regulation only affected productivity, we could allow for more flexible production technologies. We omit any fixed factor because we estimate the model in long differences over a long enough time horizon to assume labor and energy inputs are freely adjustable. When factor adjustment costs are important, the model could be extended to include a fixed factor, but then the investment decision would need to be modeled explicitly. The assumption of constant returns to scale could be relaxed with little change to the estimation procedure, but with additional notation and complexity. Additionally, since all inputs are flexibly chosen, constant returns to scale is a natural assumption.

The government sets an environmental regulation that raises the cost of emissions for a subset of firms. We express the price on emissions as a proportion of the exogenous energy price, $w_{Zt} \equiv w_{Et}(e^{\mu_t^z R_{fist}} - 1)/\kappa_t$, where w_{Et} indicates the per-unit price of energy, $R_{fist} \in \{0, 1\}$ indicates whether firm f is subject to this regulation or not at time t , and $\mu_t^z > 0$ summarizes the regulation's effect on the price of energy.⁸ Energy supply is assumed to be perfectly elastic, which is consistent with a situation in which a state-owned company supplies energy at a controlled price.

The production function can then be written as:

$$Q_{fist} = (L_{fist})^{1-\gamma} (E_{fist})^\gamma \exp(\omega_{fist} + \gamma\mu_t^e R_{fist}), \quad (5)$$

where μ_t^e reflects any endogenous adjustment of the energy efficiency (and hence, emissions efficiency) made in response to the regulation. Under the Porter Hypothesis, one would expect that firms increase energy efficiency in response to the regulation, i.e. $\mu_t^e > 0$. But we need not impose this response. It is ultimately an empirical question whether or not firms adjust efficiency as a result of the regulation.

The outside good is produced with constant returns to scale in a single input – labor – under perfect competition. Designating the outside good as numeraire, the wage rate is pinned down by labor productivity in the outside good sector: $w_{Lt} = A_{0t}$, where A_{0t} indicates unit labor requirements in the outside sector at time t . Since labor is mobile across sectors, manufacturing firms hire labor at an exogenous wage rate w_{Lt} , which does not respond to the environmental regulation.⁹

Cost minimization then yields the following unit cost function:

$$c_{fist} = (1 - \gamma)^{-(1-\gamma)} \gamma^{-\gamma} (w_{Lt})^{1-\gamma} (w_{Et})^\gamma \exp[-\omega_{fist} + \gamma(\mu_t^z - \mu_t^e) R_{fist}]. \quad (6)$$

The regulation's effect on unit cost depends on the net regulation cost, $\tau_t \equiv \gamma(\mu_t^z - \mu_t^e)$. The regulation increases a regulated firm's unit cost if the energy price effect dominates the energy efficiency effect, i.e. $\mu_t^z > \mu_t^e$, but otherwise decreases its unit cost.

⁸Note that μ_t^z is a reduced form parameter whose purpose is simply to make the regulation expressible as a proportional increase in the cost of energy for regulated firms. However, μ_t^z can in principle vary with exogenous energy price changes and changes in κ_t , so a higher energy price or lower emissions intensity in one period need not imply a higher emissions price.

⁹Essentially, we assume perfectly elastic labor supply to the manufacturing sector. We could allow for endogenous labor supply, but then we would need to estimate structural parameters of the labor market, which is beyond the scope of this paper. Additionally, we are not aware of any evidence that the EU ETS affected the wage schedules offered by firms, so we believe it is a benign assumption in the present context.

2.3 Profit Maximization and Equilibrium

We assume there exists a finite set of firms operating in each industry each year, Ω_{ist} . The set of active firms can vary over time, as exogenous industrial dynamics lead some incumbent firms to exit and some new entrants to appear, but we assume no fixed costs of operating, and do not impose zero expected profit. Hence, entry and exit are exogenous in the model.¹⁰

Each active firm in industry i at time t then chooses a price that maximizes its profits Π_{fist} , given productivity, input prices, and regulatory status. Its objective function is:

$$\max_{p_{fist}} \Pi_{fist} \equiv p_{fist} Q_{fist} - c_{fist} Q_{fist} = \left[(p_{fist})^{\frac{\rho}{\rho-1}} - c_{fist} (p_{fist})^{\frac{1}{\rho-1}} \right] (P_{ist})^{\frac{\nu-\rho}{(\nu-1)(1-\rho)}} (\Psi_{st})^{\frac{\nu}{1-\nu}} Y_{st}. \quad (7)$$

We assume that firms are monopolistic competitors, and therefore solve this maximization program without regard for how their decision affects the price index.¹¹ Due to constant returns to scale in production, the equilibrium prices of each variety f depends only on the firm's production costs and the standard Dixit-Stiglitz mark-up:

$$p_{fist} = \frac{c_{fist}}{\rho}. \quad (8)$$

Finally, to close the model, we assume that consumer expenditures on sector s at time t is equal to the sum of sales, and that national income is exogenous, as in Helpman et al. (2008).

¹⁰The assumption of exogenous entry and exit is not required for the estimation of reduced-form parameters, but it is necessary for solving counterfactuals with discrete firms, absent estimates of fixed costs of production. Models with endogenous entry and exit usually aggregate over a continuum of firms and perform counterfactuals at the industry level, in which case individual firms play no role, as in for example, Shapiro & Walker (2018). We deviate from this approach because we study a discrete set of firms that are observed over time. We show in auxiliary regressions, reported in Appendix Figure A.1, that entry and exit do not generally correlate with industry-wide regulation density in our empirical context.

¹¹Models of monopolistic competition usually feature a continuum of firms, so that market shares are infinitesimally small by construction. With a finite set of firms, though, a large firms' price decision may have a non-trivial effect on the industry-wide price indices. We maintain the monopolistic competition assumption because we believe it is a better description of the behavior of the overwhelming majority of firms, and because it simplifies the analysis quite a lot. In Appendix C.2, we consider the case of Bertrand-Nash pricing, as in Atkeson & Burstein (2008), and use this alternative assumption to perform an adversarial test of our estimation strategy. Even in that setting, our estimator reliably outperforms the alternatives.

2.4 The Effect of Regulation on Revenues

Substituting equations (6) and (8) into (2), we see that the revenues of any given firm depends on the firm's own regulation status – through the firm's own unit cost and price – and the regulation status of all other firms in the sector – through the industry and sector price indices. Taking the log change of revenues over time between pre-regulation period t_0 and post-regulation period t , we have

$$\begin{aligned} \Delta y_{fist} &= \frac{\rho\tau_t}{\rho-1}R_{fist} + \frac{\nu-\rho}{(\nu-1)(1-\rho)}\Delta P_{ist} + \frac{\nu}{1-\nu}\Delta\Psi_{st} \\ &+ \frac{\rho(1-\gamma)}{\rho-1}\Delta w_{Lt} + \frac{\rho\gamma}{\rho-1}\Delta w_{Et} + \Delta Y_{st} + \frac{\rho}{1-\rho}\Delta A_{fist} \end{aligned} \quad (9)$$

where the Δ -operator denotes the log difference between t and t_0 and $\Delta A_{fist} \equiv \omega_{fist} - \omega_{fist_0}$.

If ΔP_{ist} and $\Delta\Psi_{st}$ were observed, equation (9) could be estimated by OLS, which would identify structural parameters τ_t , ρ , and ν .¹² However, these parameters alone would not be enough to identify the treatment effects of the regulation, since the price indices themselves would be different in the counterfactual unregulated equilibrium. Thus, to evaluate counterfactuals and compute treatment effects, we also need to know how ΔP_{ist} and $\Delta\Psi_{st}$ depend on the regulation.¹³

Leveraging results from price index theory, we can express the log change in the price indices between year t_0 and t as a weighted average of log changes in individual firm prices:

$$\Delta P_{ist} = \sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta p_{kist} + \frac{1-\rho}{\rho} \Delta \lambda_{ist} \quad (10)$$

$$\Delta \Psi_{st} = \sum_{m \in \Upsilon_s} \Phi_{mst} \Delta P_{mst} \quad (11)$$

where Ω_{ist}^* denotes the set of varieties from industry i that are sold in both t_0 and t (also referred to as the continuing good set), and $\Delta \lambda_{it}$ indicates the log change in the market share of the continuing good set between t_0 and t , with $\lambda_{ist} \equiv \left(\sum_{\ell \in \Omega_{ist}^*} y_{list} \right) / \left(\sum_{\ell \in \Omega_{ist}} y_{list} \right)$ and $\lambda_{ist_0} \equiv \left(\sum_{\ell \in \Omega_{ist}^*} y_{lsit_0} \right) / \left(\sum_{\ell \in \Omega_{ist_0}} y_{list_0} \right)$.¹⁴ The terms ϕ_{kist} and Φ_{mst} denote the weights

¹²Endogeneity of price indices could be addressed by instrumental variables. See Costinot et al. (2016).

¹³Baier & Bergstrand (2009) make a similar point with respect to estimating counterfactual outcomes in the context of international trade. They observe that consistent estimation of trade elasticities requires controlling for multilateral resistance terms, which are similar to the industry-wide price indices in our context. They also argue that, in a counterfactual scenario in which trade costs change, the multilateral resistance terms would update as well, and hence need to be recomputed for the counterfactual scenario.

¹⁴Appendix B.1 derives these expressions for the changes in CES price indices, following Feenstra (1994).

applied to the price changes for firm k and industry m in sector s . As shown by Sato (1976) and Vartia (1976), these weights can be solved for analytically as functions of market shares in both periods t_0 and t .¹⁵

Substituting expressions for price-index changes (10) and (11) into equation (9) yields

$$\begin{aligned}
\Delta y_{fist} &= \frac{\rho\tau_t}{\rho-1}R_{fist} + \frac{(\nu-\rho)\tau_t}{(\nu-1)(1-\rho)} \sum_{k \in \Omega_{ist}^*} \phi_{kist}R_{kist} + \frac{(\nu-\rho)}{(\nu-1)(\rho-1)} \sum_{k \in \Omega_{ist}^*} \phi_{kist}\Delta A_{kist} \\
&+ \frac{\nu\tau}{1-\nu} \sum_{m \in \Upsilon_s} \Phi_{mst} \sum_{\ell \in \Omega_{mst}^*} \phi_{\ell mst}R_{\ell mst} + \frac{\nu}{\nu-1} \sum_{m \in \Upsilon_s} \Phi_{mst} \sum_{\ell \in \Omega_{mst}^*} \phi_{\ell mst}\Delta A_{\ell mst} \\
&+ \frac{(\nu-\rho)}{\rho(\nu-1)}\Delta\lambda_{ist} + \frac{\nu(1-\rho)}{\rho(1-\nu)} \sum_{m \in \Upsilon_s} \Phi_{mst}\Delta\lambda_{mst} + \Delta Y_{st} + \frac{\rho}{1-\rho}\Delta A_{fist}. \tag{12}
\end{aligned}$$

This equation will serve as the basis for estimation in Section 3. The first term on the right hand side captures the direct effect of the regulation on revenues (i.e. holding price indices constant). This term can be positive or negative, depending on the net regulation cost, τ_t . The second and third terms capture the regulation's indirect effect through the price index of the continuing good set within an industry. The second term depends on the combined market share of regulated firms within the industry. The larger the market share of regulated firms, the larger the indirect effect.¹⁶ The third term depends on the regulation through the Sato-Vartia weights, ϕ_{kist} , associated to firm-level productivity shocks. The fourth and fifth terms are just sector-level analogues of the previous two terms. The last line of equation (12) captures the effect on revenues of entry and exit, exogenous movements in aggregate expenditures, and the firm's own idiosyncratic productivity shock.¹⁷

¹⁵Following Sato (1976) and Vartia (1976), we can solve for the weights as follows:

$$\begin{aligned}
\phi_{kist} &\equiv \frac{\vartheta_{kist} - \vartheta_{kist_0}}{\ln \vartheta_{kist} - \ln \vartheta_{kist_0}} \bigg/ \sum_{\ell \in \Omega_{ist}^*} \frac{\vartheta_{\ell ist} - \vartheta_{\ell ist_0}}{\ln \vartheta_{\ell ist} - \ln \vartheta_{\ell ist_0}} \\
\Phi_{mst} &\equiv \frac{\Theta_{mst} - \Theta_{mst_0}}{\ln \Theta_{mst} - \ln \Theta_{mst_0}} \bigg/ \sum_{h \in \Upsilon_s} \frac{\Theta_{hst} - \Theta_{hst_0}}{\ln \Theta_{hst} - \ln \Theta_{hst_0}}
\end{aligned}$$

with term $\vartheta_{kist} \equiv y_{kist} / \sum_{\ell \in \Omega_{ist}^*} y_{\ell ist}$ denotes the market share of firm k in the continuing good set of all varieties from industry i at time t , and Θ_{mst} denotes the market share of industry m in sector s at time t (assuming that the set of industries does not change over time).

¹⁶To the extent that other types of spillover effects also break down along industrial categorization, our approach will pick up some of these effects too. For example, if technological adoption at one firm in the industry encourages other firms in the same industry to adopt, these effects will to some degree be reflected in our industry-wide treatment density terms.

¹⁷Input prices w_{Lt} and w_{Et} drop out of the equation, by virtue of being common to all firms. Take any component of costs x , like input prices, that are the same for all firms. This component will cancel out of

2.5 The Effect of Regulation on Emissions

The first order conditions for cost minimization yields the following expression for firm f 's emissions at time t :

$$Z_{fist} = \rho\gamma\kappa_t y_{fist} (w_{Et} e^{\mu_t^z R_{fist}})^{-1}. \quad (13)$$

From this equation, we see that the emission intensity of revenue (Z_{fist}/y_{fist}) is constant within the sector for firms with the same regulatory status, R_{fist} . However, the emission intensity of output (Z_{fist}/Q_{fist}) is declining in productivity. Hence, more productive firms tend to emit less pollution per unit of physical output.¹⁸

The change in emissions from the base year t_0 to period t can thus be written as:

$$\Delta Z_{fist} = \Delta\kappa_t + \Delta y_{fist} - \Delta w_{Et} - \mu_t^z R_{fist}, \quad (14)$$

where Δy_{fist} is given by equation (12). From this expression, we see that the regulation affects the firm's emissions both by altering the scale of economic output, Δy_{fist} , and by encouraging substitution away from the polluting input, through $-\mu_t^z$. We will refer to these as the "scale" and "technique" effects, respectively. The terms $\Delta\kappa_t$ and Δw_{Et} are not affected by the regulation, and thus adjust exogenously.

2.6 Extension to Multi-plant Firms

We now extend our theoretical framework to allow for multi-plant firms. There are three reasons this extension is important in our context. First, the regulation we study distinguishes between plants within the firm. Hence, it is necessary to map plant-level regulation indicators to firm-level sales outcomes, which requires some assumption on how multi-plant firms behave. Second, the emissions data are observed at the plant level, and only for a subset of plants. Hence, we need to model the emissions outcomes of individual plants, not firms. Finally, identification in the EU ETS literature is often based on variation in the internal distribution of economic activity across firms, which inherently relies on a multi-plant structure. With this extension, we specify microfoundations that make it possible to

equation (12) since

$$\frac{\rho x}{1-\rho} + \frac{(\nu-\rho) \sum_{k \in \Omega_{ist}^*} \phi_{kist} x}{(\nu-1)(\rho-1)} + \frac{\nu \sum_{m \in \Upsilon_s} \Phi_{mst} \sum_{\ell \in \Omega_{mst}^*} \phi_{\ell mst} x}{\nu-1} = x \left[\frac{\rho}{1-\rho} + \frac{(\nu-\rho)}{(\nu-1)(\rho-1)} + \frac{\nu}{\nu-1} \right] = 0$$

¹⁸Evidence of the negative relationship between economic productivity and emission intensity can be found in many empirical contexts (Holladay, 2016; Forslid et al., 2018; Barrows & Ollivier, 2018).

assess the assumptions implicit in this standard identification strategy.

In the multi-plant version of the model, firm f still produces a unique variety f , but now owns J_{fist} plants at time t . These plants produce intermediate outputs using Cobb-Douglas combinations of labor and energy, and the intermediate outputs are then aggregated into a final output.¹⁹ We assume that the Hicks-neutral productivity, ω_{fist} , is common to all plants $j \in \{1, 2, \dots, J_{fist}\}$ within the firm. Regulation, however, can be plant-specific, with $R_{jfirst} = 1$ indicating that plant j is subject to regulation at time t , and zero otherwise. Firm f 's output is then given by:

$$Q_{fist} = \prod_{j=1}^{J_{fist}} \left(\alpha_{jfirst}^{-\alpha_{jfirst}} \right)^{1-\sigma} \left[(L_{jfirst})^{1-\gamma} (E_{jfirst})^\gamma \exp(\omega_{fist} + \gamma\mu^e R_{jfirst}) \right]^{\alpha_{jfirst}}, \quad (15)$$

where α_{jfirst} denotes the contribution of plant j to firm f 's total output, with $\sum_j \alpha_{jfirst} = 1$. The parameter σ captures the relationship between overall production and the degree of dispersion across plants. If dispersing economic activity mainly increases cost and reduces production, then $\sigma > 0$. On the other hand, if dispersing economic activity creates sufficiently large economies of scope that reduce the unit cost and increase production, then $\sigma < 0$.

Cost minimization yields the following firm-level unit cost function:

$$c_{fist} = (1-\gamma)^{-(1-\gamma)} \gamma^{-\gamma} \left(\prod_{j=1}^{J_{fist}} \alpha_{jfirst}^{-\alpha_{jfirst}} \right)^\sigma (w_{Lt})^{1-\gamma} (w_{Et})^\gamma \exp \left(-\omega_{fist} - \gamma(\mu^e - \mu^z) \sum_j \alpha_{jfirst} R_{jfirst} \right). \quad (16)$$

This expression is almost identical to equation (6), except that we now account for variation in regulatory status across each firm's plants and that we include a term to capture the relationship between unit cost and the dispersion of economic activity across plants, $\prod_{j=1}^{J_{fist}} (\alpha_{jfirst}^{-\alpha_{jfirst}})^\sigma$. Cost minimization implies $L_{jfirst} = \alpha_{jfirst} L_{fist}$, thereby allowing us to compute firm-level treatment density $R_{fist} \equiv \sum_j \alpha_{jfirst} R_{jfirst} = \sum_j \frac{L_{jfirst}}{L_{fist}} R_{jfirst}$ and dispersion measure $\alpha_{fist} \equiv \prod_{j=1}^{J_{fist}} \alpha_{jfirst}^{-\alpha_{jfirst}} = \prod_{j=1}^{J_{fist}} \left(\frac{L_{jfirst}}{L_{fist}} \right)^{-\left(\frac{L_{jfirst}}{L_{fist}} \right)}$ as functions of plant-level

¹⁹This mapping of intermediate outputs to the final product is reminiscent of the approach of aggregating tasks to produce a final good, as in Acemoglu & Autor (2011). Presumably, firms own multiple plants because there exists complementarities across plants, which could stem from economies from vertical integration that improve productivity and lower overall costs of production (Alfaro et al., 2016), or from complementarities in intangible inputs, such as high-quality management, marketing know-how, or R&D capital (Atalay et al., 2014).

labor shares.²⁰

Collecting structural parameters into reduced-form parameters, the changes in firm-level revenues between a pre-regulation period t_0 and a post-regulation period t can be written as

$$\begin{aligned}
\Delta y_{fist} &= \beta_{0,t} R_{fist} + \beta_1 \Delta \alpha_{fist} + \beta_{2,t} \sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist} + \beta_3 \sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta \alpha_{kist} + \beta_4 \Delta \lambda_{ist} \\
&+ \beta_{5,t} \sum_{m \in \Upsilon_s} \Phi_{mst} \sum_{\ell \in \Omega_{mst}^*} \phi_{\ell mst} R_{\ell mst} + \beta_6 \sum_{m \in \Upsilon_s} \Phi_{mst} \sum_{\ell \in \Omega_{mst}^*} \phi_{\ell mst} \Delta \alpha_{\ell mst} + \beta_7 \sum_{m \in \Upsilon_s} \Phi_{mst} \Delta \lambda_{mst} \\
&+ \beta_8 \sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta A_{kist} + \beta_9 \sum_{m \in \Upsilon_s} \Phi_{mst} \sum_{\ell \in \Omega_{mst}^*} \phi_{\ell mst} \Delta A_{\ell mst} + \Delta Y_{st} + \xi_{fist}, \tag{17}
\end{aligned}$$

with $\xi_{fist} = \frac{\rho}{1-\rho} \Delta A_{fist}$.²¹ Compared to equation (12), our multi-plant extension adds three terms: the exogenous change in firm f 's dispersion of activities, and their industry- and sector-level counterparts.

Moving to the environmental effect of the regulation, cost minimization at the plant level implies

$$\Delta Z_{jfst} = -\mu_t^z R_{jfst} + \Delta \alpha_{jfst} + \Delta \kappa_t - \Delta w_{Et} + \Delta y_{fist}. \tag{18}$$

The changes in plant-level emissions between pre-period t_0 and period t are thus affected by a firm-level scale effect, Δy_{fist} , a plant-level technique effect, $-\mu_t^z$, and a change in plant-specific labor share, $\Delta \alpha_{jfst}$.

2.7 Average and Aggregate Treatment Effects

As long as a regulation affects energy prices ($\mu_t^z \neq 0$) and the unit cost of production ($\tau_t \neq 0$), it will influence the revenues and emissions of both regulated and unregulated firms. The magnitude of treatment effects vary across firms, though, due to differences in the density of regulated firms across industries. With multi-plant firms, treatment effects also vary because of differences in labor shares of treated plants within regulated firms. Our main objective is to estimate the averages and aggregates of these firm-level effects.

²⁰Although fixed capital does not enter explicitly in our model, these plant-level labor shares are effectively a proxy for capital shares—the fixed capital at each plant determines the optimal allocation of labor across plants.

²¹Other parameters are defined as $\beta_{0,t} \equiv \frac{\rho \tau_t}{\rho-1}$, $\beta_1 \equiv \frac{\rho \sigma}{\rho-1}$, $\beta_{2,t} \equiv \frac{(\nu-\rho)\tau_t}{(\nu-1)(1-\rho)}$, $\beta_4 = \frac{(\nu-\rho)}{\rho(\nu-1)}$, $\beta_3 \equiv \frac{(\nu-\rho)\sigma}{(\nu-1)(1-\rho)}$, $\beta_{5,t} \equiv \frac{\nu \tau_t}{1-\nu}$, $\beta_6 \equiv \frac{\nu \sigma}{1-\nu}$, $\beta_7 \equiv \frac{\nu(1-\rho)}{\rho(1-\nu)}$, $\beta_8 \equiv \frac{(\nu-\rho)}{(\nu-1)(\rho-1)}$, $\beta_9 \equiv \frac{\nu}{\nu-1}$.

At the firm level, we define the treatment effect for outcome $v \in \{y, Z\}$ as the log difference between the observed post-regulation outcome and the unregulated counterfactual outcome: $\ln\left(\frac{v_{fist}}{v'_{fist}}\right) = \Delta v_{fist} - \Delta v'_{fist}$, where v'_{fist} indicates the estimated outcome for firm f in year t in the case that no firm was regulated. We then define the average treatment effect for treated (T) and control (C) firms, respectively, with $X \in \{T, C\}$, as follows:

$$ATX_t^v \equiv \frac{1}{N_t^X} \sum_{f \in \Omega_t^X} \ln\left(\frac{v_{fist}}{v'_{fist}}\right) = \frac{1}{N_t^X} \sum_{f \in \Omega_t^X} (\Delta v_{fist} - \Delta v'_{fist}) \quad (19)$$

where N_t^X and Ω_t^X denotes the total number of firms and the set of firms in group X at time t , respectively.

At the plant level, we distinguish between three groups: regulated plants (denoted TT , for treated plants in treated firms), unregulated plants owned by firms that also operate regulated plants (CT), and unregulated plants owned by firms without any regulated plant (CC). Then, for any group of plants, $X \in \{TT, CT, CC\}$, we define the average treatment effect on emissions as:

$$ATX_t^Z = \frac{1}{N_t^X} \sum_{j \in \Omega_t^X} \ln\left(\frac{Z_{jfist}}{Z'_{jfist}}\right), \quad (20)$$

where N_t^X and Ω_t^X denote the number and set of plants in year t for group X .

We are also interested in estimating the effect on aggregate emissions.²² We define the aggregate effect as:

$$ATZ_t = \ln\left(\frac{\sum_f Z_{fist}}{\sum_f Z'_{fist}}\right). \quad (21)$$

3 Estimation and Monte Carlo Evidence

In this section, we present our strategy for estimating the effects of regulation on firm-level revenues and plant-level emissions based on equations (17) and (18), alongside commonly used alternative strategies. Finally, we evaluate the finite sample properties of these estimators in Monte Carlo experiments.

²²There are no effects of the regulation on aggregate manufacturing sales by assumption.

3.1 Our GMM Procedure

Estimating Reduced-Form Parameters. The model implies that the change in firm-level revenues between a pre-regulation period and a post-regulation period can be expressed as a linear function of regulation assignments, market shares, and unobserved productivity shocks. Collecting all sector-wide variables from equation (17) into sector-time fixed effects, δ_{st} , yields the following simplified expression:²³

$$\begin{aligned} \Delta y_{fist} &= \beta_0 R_{fist} + \beta_1 \Delta \alpha_{fist} + \beta_2 \sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist} + \beta_3 \sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta \alpha_{kist} + \beta_4 \Delta \lambda_{ist} \\ &+ \beta_8 \sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta A_{kist} + \delta_{st} + \frac{\rho}{1 - \rho} \Delta A_{fist}. \end{aligned} \quad (22)$$

All variables in the first line of (22) are observed, while variables in the second line are not.

If productivity follows a random walk, then ΔA_{fist} would be uncorrelated with the regulation and with pre-period characteristics. In this case, it would be tempting to group ΔA_{fist} and $\sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta A_{kist}$ into the error term and estimate (22) by OLS. However, even with random-walk productivity growth, two identification problems remain. First, the overall market share of regulated firms within industries, $\sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist}$, is correlated with the unobserved weighted sum of productivity shocks, $\sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta A_{kist}$. To see this, imagine an industry with no regulated firms. In this industry, the market shares ϕ_{kist} depend exclusively on firms' productivities, A_s , and on the dispersion of their economic activities, α_s . But, as regulation density in the industry increases, the market shares will increasingly depend on regulation, and less on A_s and α_s . Hence, industries with higher regulation density (high $\sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist}$) will have lower correlations between ΔA_{kist} and ϕ_{kist} , and hence low values of $\sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta A_{kist}$.²⁴

The second identification problem comes from the fact that industries with higher weighted-average productivity growth of continuing firms (high values of $\sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta A_{kist}$) will mechanically have higher revenue share growth of continuing firms. Hence, we would expect $\Delta \lambda_{ist}$ to correlate with $\sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta A_{kist}$ as well.

²³For simplicity, we drop the time index from parameters β_s , which originates from the time-varying τ_t , in the estimation equation. Yet, we will estimate parameters in long differences for each period, so these parameters are allowed to vary over time.

²⁴This relationship is likely non-linear, as the market shares in an industry with all regulated firms would also depend purely on A_s and α_s . Note that this endogeneity problem cannot be solved simply by instrumenting $\sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist}$ with base-year weighted average regulation. This is because base-year weighted average regulation would correlate with $\sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta A_{kist}$ as well. Rather, we would need an instrument that does not depend on the regulation.

To address these identification problems, we build a GMM estimator to simultaneously estimate β_0, β_1 , and β_2 . The first step is to re-write (22) as a system of equations:

$$\begin{aligned}\Delta y_{fist} &= \beta_0 R_{fist} + \beta_1 \Delta \alpha_{fist} + \eta_{ist} + \xi_{fist} \\ \eta_{ist} &= \beta_2 \sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist} + \beta_1 \frac{\beta_2}{\beta_0} \sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta \alpha_{kist} - \frac{\beta_2}{\beta_0} \left(\Delta \lambda_{ist} - \sum_{k \in \Omega_{ist}^*} \phi_{kist} \xi_{fist} \right) + \delta_{st} + \varsigma_{ist},\end{aligned}$$

with $\xi_{fist} \equiv \frac{\rho}{1-\rho} (A_{fist} - \overline{\Delta A_{ist}})$ and $\varsigma_{ist} \equiv \frac{\nu}{1-\nu} (\overline{\Delta A_{ist}} - \overline{\Delta A_{st}})$, and where we exploit the following relationships: $\beta_3 = \beta_1 \beta_2 / \beta_0$, $\beta_4 = -\beta_2 / \beta_0$, and $\beta_8 = -(\beta_2 / \beta_0) \rho / (1 - \rho)$.²⁵ Notice that the first of these equations is specified at the firm level, and the second one is specified at the industry level. Then, for any candidate vector $\beta^* = (\beta_0^*, \beta_1^*, \beta_2^*)$, we can compute the following residuals:

$$\widehat{\xi}_{fist} = \Delta y_{fist} - \beta_0^* R_{fist} - \beta_1^* \Delta \alpha_{fist} - \left[\frac{1}{N_{ist}^*} \sum_{k \in \Omega_{ist}^*} \Delta y_{kist} - \beta_0^* R_{kist} - \beta_1^* \Delta \alpha_{kist} \right] \quad (23)$$

and

$$\begin{aligned}\widehat{\varsigma}_{ist} &= \left[\frac{1}{N_{ist}^*} \sum_{k \in \Omega_{ist}^*} \Delta y_{kist} - \beta_0^* R_{kist} - \beta_1^* \Delta \alpha_{kist} \right] - \beta_2^* \sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist} \\ &\quad - \beta_1^* \frac{\beta_2^*}{\beta_0^*} \sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta \alpha_{kist} + \frac{\beta_2^*}{\beta_0^*} \left(\Delta \lambda_{ist} - \sum_{k \in \Omega_{ist}^*} \phi_{kist} \widehat{\xi}_{fist} \right) - \delta_{st},\end{aligned} \quad (24)$$

where N_{ist}^* represents the number of continuing firms in industry i in sector s in post-regulation period t .²⁶

To evaluate a candidate vector β^* , we exploit the following moment conditions: (1) $E[\xi_{fist} R_{fist}] = 0$, (2) $E[\xi_{fist} \Delta \alpha_{fist}] = 0$, (3) $E[\varsigma_{ist} (\sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist})] = 0$, and (4) $E[\varsigma_{ist} (\sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta \alpha_{kist})] = 0$. The first and third moment conditions state that the firm-level and industry-level residuals, respectively, are uncorrelated with the regulation coverage at the corresponding levels. Both of these moment conditions will hold auto-

²⁵Other parameters are also functions of $\{\beta_0, \beta_1, \beta_2\}$ since $\beta_5 = -(\beta_0 + \beta_2)$, $\beta_6 = -(\beta_1 + \beta_3)$, $\beta_7 = (\beta_0 + \beta_2) / \beta_0$, and $\beta_9 = (1 + \beta_2 / \beta_0) \rho / (1 - \rho)$.

²⁶To control for the sector-year fixed effects, we first compute

$$\left[\frac{1}{N_{ist}^*} \sum_{k \in \Omega_{ist}^*} \Delta y_{kist} - \beta_0^* R_{kist} - \beta_1^* \Delta \alpha_{kist} \right] - \beta_2^* \sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist} - \beta_1^* \frac{\beta_2^*}{\beta_0^*} \sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta \alpha_{kist} + \frac{\beta_2^*}{\beta_0^*} \left(\Delta \lambda_{ist} - \sum_{k \in \Omega_{ist}^*} \phi_{kist} \widehat{\xi}_{fist} \right),$$

then regress this quantity on a vector of sector-by-year indicator variables and extract the residual.

matically if productivity follows a random walk, since past productivity growth, even if it predicts regulation, would not predict future productivity growth. If productivity does not follow a random walk, however, future productivity growth will be correlated with pre-period characteristics, which may also predict firm-level and industry-level regulation coverage. In this case, these moment conditions would still hold for an appropriately matched sample of firms and industries (we provide details on how to operationalize this in Section 4.2). This is because, for any subset of regulated and unregulated firms with the same pre-period productivity, future productivity growth would be uncorrelated with the regulation, absent a treatment effect. The second and fourth moment conditions follow from the exogeneity of α_{fist} .

Using estimated residuals (23) and (24), we build the following empirical moments:

$$\mathbf{\Gamma}(\beta_0^*, \beta_1^*, \beta_2^*) = \begin{pmatrix} \sum_t \sum_{f,i} \widehat{\xi}_{fist}(\beta_0^*, \beta_1^*) R_{fist} \\ \sum_t \sum_{f,i} \widehat{\xi}_{fist}(\beta_0^*, \beta_1^*) \Delta\alpha_{fist} \\ \sum_t \sum_i \widehat{\varsigma}_{ist}(\beta_0^*, \beta_1^*, \beta_2^*) \left(\sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist} \right) \\ \sum_t \sum_i \widehat{\varsigma}_{ist}(\beta_0^*, \beta_1^*, \beta_2^*) \left(\sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta\alpha_{kist} \right) \end{pmatrix}$$

and we choose β to minimize the usual GMM criterion function:

$$\beta^{STR} = \min_{\beta_0^*, \beta_1^*, \beta_2^*} \mathbf{\Gamma}(\beta_0^*, \beta_1^*, \beta_2^*)' \mathbf{W} \mathbf{\Gamma}(\beta_0^*, \beta_1^*, \beta_2^*)$$

where the weighting matrix \mathbf{W} is estimated optimally using the two-step GMM estimator.

Finally, we estimate μ_z from plant-level emissions data. Moving Δy_{fist} and $\Delta \alpha_{jfirst}$ to the left hand side of equation (18), we have:

$$\Delta \left(\frac{Z_{jfirst}}{\alpha_{jfirst} y_{fist}} \right) = \beta_z R_{jfirst} + \delta_{ist} + \epsilon_{jfirst}, \quad (25)$$

where δ_{ist} absorbs all determinants of emissionintensity that are common across plants within an industry and ϵ_{jfirst} captures measurement error in emissions, which we assume is uncorrelated with R_{jfirst} . Hence, we have $E[\widehat{\beta}_z] = -\mu_z$, the input price effect of the regulation on emissions, where $\widehat{\beta}_z$ results from estimating 25 by OLS.

Counterfactuals and Treatment Effects. Once we have unbiased estimates of the four crucial reduced-form parameters— β_0 , β_1 , β_2 , and β_z —the final step is to compute the average and aggregate treatment effects, relative to the counterfactual equilibrium where

no firms are regulated. This calculation allows all endogenous terms to update. To see this, we can derive the counterfactual revenues at time t for a firm with positive revenues in the post-regulation period using equation (17):

$$\begin{aligned}
\ln(y'_{fist}) &= \ln(y_{fist}) - \widehat{\beta}_0 R_{fist} - \widehat{\beta}_2 \sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist} + (\widehat{\beta}_0 + \widehat{\beta}_2) \sum_{m \in \Upsilon_s} \Phi_{mst} \sum_{\ell \in \Omega_{mst}^*} \phi_{\ell mst} R_{\ell mst} \\
&+ \frac{\widehat{\beta}_2}{\widehat{\beta}_0} \sum_{k \in \Omega_{ist}^*} (\phi'_{kist} - \phi_{kist}) \widehat{\xi}_{kist} - \frac{\widehat{\beta}_2}{\widehat{\beta}_0} (\Delta \lambda'_{ist} - \Delta \lambda_{ist}) + \widehat{\beta}_1 \frac{\widehat{\beta}_2}{\widehat{\beta}_0} \sum_{k \in \Omega_{ist}^*} (\phi'_{kist} - \phi_{kist}) \Delta \alpha_{kist} \\
&+ \frac{\widehat{\beta}_0 + \widehat{\beta}_2}{\widehat{\beta}_0} \left(\sum_{m \in \Upsilon_s} \Phi'_{mst} \Delta \lambda'_{mst} - \sum_{m \in \Upsilon_s} \Phi_{mst} \Delta \lambda_{mst} \right) - \sum_{m \in \Upsilon_s} (\Phi'_{mst} - \Phi_{mst}) \widehat{\varsigma}_{mst} \\
&- \frac{\widehat{\beta}_0 + \widehat{\beta}_2}{\widehat{\beta}_0} \sum_{m \in \Upsilon_s} \sum_{\ell \in \Omega_{mst}^*} (\Phi'_{mst} \phi'_{\ell mst} - \Phi_{mst} \phi_{\ell mst}) \left(\widehat{\xi}_{\ell mst} + \widehat{\beta}_1 \Delta \alpha_{\ell mst} \right), \tag{26}
\end{aligned}$$

where ϕ'_{fist} , Φ'_{ist} and $\Delta \lambda'_{ist}$ respectively indicate counterfactual Sato-Vartia weights and continuing-good share growth, which must be solved for endogenously. Terms $\widehat{\varsigma}_{mst}$ and $\widehat{\xi}_{kist}$ are computed according to (23) and (24), whereas $\widehat{\beta}_0$, $\widehat{\beta}_1$ and $\widehat{\beta}_2$ result from the GMM procedure.

Since equation (26) holds for each firm f , this system of non-linear equations can be solved for the vector of counterfactual revenues. The equilibrium can be found with either a numerical solver or a fixed-point algorithm.²⁷ Once we have estimates of y'_{fist} , counterfactual emissions at the plant level are easily computed as $Z'_{j fist} = Z_{j fist} e^{-\widehat{\beta}_2 R_{j fist} (y'_{fist}/y_{fist})}$, which can be aggregated at the firm level, $Z'_{fist} = \sum_j Z'_{j fist}$.

From the vectors of actual and counterfactual revenues and emissions, it is straightforward to compute treatment effects at the firm level, plant level, and aggregate level following the expressions (19), (20), and (21), respectively.

3.2 Evaluating Treatment Effects With Existing Estimators

Previous literature typically estimates the effects of incomplete regulation using either the difference-in-differences estimator (DD), or what we refer to as the local approximation estimator (LA) (see for example Rotemberg 2019; Muehlegger & Sweeney 2021; Cai & Szeidl 2022). Both approaches are valid under special cases of our data generating process, but neither estimator delivers unbiased estimates of treatment effects in general.

²⁷In the fixed point algorithm, we first set $y'_{fist} = y_{fist}$ for each f , compute ϕ'_{fist} , Φ'_{ist} and λ'_{ist} under this assumption, compute the right hand side of (26), and then update y'_{fist} . The process then repeats until the vector $\{y'_{fist}\}$ converges.

Difference-in-Differences Estimation. In the multi-plant model, we can collect variables into industry-year fixed effects and express the changes in outcome $v \in \{y, Z\}$ between pre-regulation and a post-regulation periods as:

$$\Delta v_{fist} = \beta_0^{v,DD} R_{fist} + \beta_1^{v,DD} \Delta \alpha_{fist} + \delta_{ist} + \epsilon_{fist} \quad (27)$$

where $R_{fist} \equiv \sum_j \alpha_{jfist} R_{jfist}$.²⁸ This expression resembles a conventional two-way fixed effect model with continuous treatment (expressed in long differences), with an additional control for the dispersion of economic activity within firms. According to the assumed data generating process, δ_{ist} absorbs all effects of the regulation on the industry and sector-wide price indices, as well as all other industry-wide effects, and the error term is $\epsilon_{fist} \equiv \frac{\rho}{1-\rho} \Delta A_{fist}$. Hence, if regulation is orthogonal to productivity growth, conditional on industry-year, we have $E \left[\widehat{\beta_0^{y,DD}} \right] = \frac{\rho\tau}{\rho-1}$ and $E \left[\widehat{\beta_0^{z,DD}} \right] = -\mu^z + \frac{\rho\sigma}{\tau-1}$. Researchers often rely on matched samples to justify the assumption that the regulation is orthogonal to firms' productivity growth. When this condition holds, the DD estimator delivers unbiased estimates of the direct effect of the regulation on firm-level revenues and emissions.²⁹

Nevertheless, even under conditions where $\beta_0^{v,DD}$ is an unbiased estimator of the direct effect of the regulation, β_0 , this estimate ought not be interpreted as an average treatment effect. If cost shocks are passed through to firm-level prices, the regulation will affect price indices, and the direct effect of regulation will diverge from the average treatment effect, ATT_t^v .

When each firm operate a single plant, we can characterize the discrepancy between the DD estimator and the true ATT_t^v analytically. With single-plant firms, R_{fist} becomes binary and $\Delta \alpha_{fist}$ drops out because it equals zero for all firms. In this case, we can show that the $E \left[\widehat{\beta_0^{y,DD}} \right]$ is greater in magnitude than the true ATT_t^y (see propositions 1 and 2 in Appendix B.3). For firm-level emissions, by contrast, we generally cannot sign the difference between the DD estimator and the true ATT_t^Z .

These analytical results are depicted in Figure 1. The solid black line in Figure 1 (left

²⁸To allow regression coefficients to vary by year, we suppose that this regression is run in long differences year by year. Our specification in long differences eliminates any bias that would be generated by a correlation between time-invariant firm characteristics and the regulation status.

²⁹In practice, researchers typically simplify equation (27) by representing regulation with a binary indicator for whether the firm operates any regulated plants, i.e. $R_{fist} \equiv \max\{R_{jfist}\}$, and ignoring the dispersion of economic activity within firms, as captured by $\Delta \alpha_{fist}$. When treatment is not randomly assigned to plants, the dispersion measure may correlate with treatment, which generates omitted variable bias, according to our multi-plant model.


panel) plots the true ATT for revenues as a function of μ_t^e . When the energy price effect, μ_t^z , is larger than the energy efficiency effect, μ_t^e , the ATT for revenues is negative, and vice versa when $\mu_t^e > \mu_t^z$. The dashed black line, which represents the DD estimate of the ATT_t^y , passes through zero for $\mu_t^e = \mu_t^z$, but has a steeper slope than the solid line because the DD estimate of ATT_t^y ignores the spillovers (propositions 1 and 2). Hence, the DD estimator exaggerates the magnitude of the true ATT_t^y .

The relationship between the true ATT_t^Z and the DD estimate of ATT_t^Z is not as straightforward. The true ATT and the DD estimator for emissions, depicted by the solid and dashed blue lines, respectively, have the same relationship to μ_t^e as do their counterparts for revenues, but are shifted downwards by μ_t^z . Even at the point where $\mu_t^e = \mu_t^z$, where there is no scale effect, the true ATT_t^Z is negative because firms substitute away from energy ($\mu_t^z > 0$). As μ_t^e gets larger, the scale effect increases, even to the point where it can dominate (right of C), in which case the ATT for emissions will be positive. This creates three distinct possibilities. For points to the left of A or right of C, the DD estimator will exaggerate the magnitude of the true ATT_t^Z . For points between A and B, the DD estimator will understate the true ATT_t^Z . And for intermediate values of μ_t^e , between B and C, the DD estimator would have the opposite sign of the true ATT_t^Z . It is because of this third possibility that we cannot infer the sign of the treatment effect from the DD estimator.

To translate the DD estimate into an effect on aggregate emissions, we compute counterfactuals by subtracting out the DD coefficient from the emissions of regulated firms and sum over all firms (dashed blue line, right panel). The true ATZ_t (solid blue line, right panel) remains negative even when the firm-level ATT_t^Z is positive (right of C). This is because, as the regulated firms become more efficient, they capture greater market share (propositions 1 and 2). The DD estimator could therefore yield the correct sign of the average treatment effect, but still get the sign wrong for the aggregate effect on emissions (right of C).

The firm-level and aggregate analytical results discussed so far do not carry over to multi-plant firms. More specifically, in multi-plant settings, according to our model, we cannot say analytically if the DD estimator of ATT_t^y , ATT_t^Z , or ATZ_t is biased up or down or even has the right sign (see Appendices B.5, B.6 and B.7). The intuition for this result is that, in a multi-plant world, the sign of the true effect of regulation can vary even across firms that own regulated plants, but the DD estimator is still biased up when $\tau_t < 0$, and biased downward otherwise. Hence, cases can arise when the true ATT is negative, but

Figure 1: True Values and DD Estimates of Firm-level ATTs and Aggregate Effects



Notes: The graph on the left shows true ATT (solid lines) and DD estimates (dashed lines) of ATT for single-plant firm-level revenues (in blue) and emissions (in black). The graph on the right shows true (solid line) aggregate effects on emissions, ATZ, and the ones estimated by DD (dashed line). The grey vertical line indicates when $\mu^e = \mu^z = 0.2$. When $\mu^e < \mu^z$, then $\tau > 0$, whereas $\tau < 0$ when $\mu^e > \mu^z$.

the DD estimator is biased up, so the DD estimator is biased towards zero, for example. We discuss these results at length in appendices B.5, B.6 and B.7.

In multi-plant settings, researchers also use DD to estimate effects of regulation on different subsets of plants – for instance, unregulated plants owned by firms that also operate regulated plants. In this situation, researchers typically estimate a plant-level DD regression, comparing these indirectly regulated plants with control plants owned by firms that own no regulated plants, just as one would normally compare directly regulated plants with unregulated controls (see for example Bartram et al. 2022; Gibson 2019; Soliman 2020). For each group of plants, $X \in \{TT, CT\}$, such plant-level DD estimators can be written as:

$$\Delta Z_{j\text{fist}} = \beta_0^{Z,DD,X} R_{j\text{fist}}^X + \delta_{ist} + \epsilon_{j\text{fist}}, \quad (28)$$

where $R_{j\text{fist}}^X$ is one of two indicator variables: (1) $R_{j\text{fist}}^X \equiv R_{j\text{fist}}$ when estimating the effect on directly regulated plants (TT), or (2) $R_{j\text{fist}}^X \equiv \max\{R_{k\text{fist}}\}_{k=1}^J$ when estimating the effect on indirectly regulated plants (CT).³⁰

The coefficient $\beta_0^{Z,DD,X}$ in equation (28) is meant to capture the effect of the regulation

³⁰We assume revenues are only observed at the firm-level, as they are in our empirical setting. So we only specify these plant-level regressions for emissions. But of course, when revenues are observed by plant, one could also estimate versions of (28) taking the change in plant-level revenues as the outcome variable.

at the plant level. However, our multi-plant data generating process indicates that firm-level regulation, R_{fist} , contributes to plant-level emissions through the scale effect (see equation 18). When one plant becomes regulated, it affects R_{fist} as well as the firm-wide unit cost, which alters the optimal scale of production across all plants. Hence, omitting firm-level regulation from (28) would lead us to expect $E[R_{jfst}\epsilon_{jfst}] \neq 0$. In appendix B.7, we show that, when firms operate multiple plants, we cannot sign the difference between the DD estimator and the true ATT or ATZ .

Estimation via Local Approximation. A handful of papers extend the DD model to take into account spillover effects across firms (Cai & Szeidl, 2022; Muehlegger & Sweeney, 2021; Rotemberg, 2019). In these papers, the conventional DD regression is augmented with controls for density of treatment within a neighborhood of each firm (where neighborhood often refers to geographical proximity and is meant to proxy for a market). There is an intuitive appeal to this approach: if regulation increases costs and prices, then firms should benefit from having a high share of regulated competitors.

In our context, where competitors are defined as firms in the same industry, this augmented DD regression would take the form:

$$\begin{aligned} \Delta v_{fist} = & \beta_0^{v,LA} R_{fist} + \beta_1^{v,LA} \Delta \alpha_{fist} + \beta_2^{v,LA} \sum_{k \in \Omega_{ist_0}} \theta_{kist_0} R_{kist} \\ & + \beta_3^{v,LA} \sum_{k \in \Omega_{ist_0}} \theta_{kist_0} \Delta \alpha_{kist} + \delta_{st} + \epsilon_{fist}. \end{aligned} \quad (29)$$

where δ_{st} represents a sector-year fixed effect and θ_{kist_0} denotes the market share of firm k in industry i in pre-regulation year t_0 . In a simplified setting with single-plant firms, the $\Delta \alpha$ -terms would drop out. The spillover effect $\beta_2^{v,LA}$ is identified from variation in regulation density across industries. Since the spillover term only involves pre-period weights, which are exogenous to the regulation under a random walk assumption (or in a matched sample), we could estimate (29) via OLS, and compute firm-level treatment effects as follows:³¹

$$\ln \left(\frac{v_{fist}}{v_{LAfist}} \right) = \widehat{\beta_0^{v,LA}} R_{fist} + \widehat{\beta_2^{y,LA}} \sum_{k \in \Omega_{ist_0}} \theta_{kist_0} R_{kist} - (\widehat{\beta_0^{v,LA}} + \widehat{\beta_2^{y,LA}}) \sum_{m \in \Upsilon} \Theta_{mt_0} \sum_{k \in \Omega_{ist_0}} \theta_{kist_0} R_{kist} \quad (30)$$

We refer to (29)-(30) as the ‘‘local approximation’’ (LA) method because these equations can be derived by taking a local approximation of the change in the price index around

³¹The third term in the expression of $\ln \left(\frac{v_{fist}}{v_{LAfist}} \right)$ follows from the fact that $\frac{\nu \tau_t}{1-\nu} = - \left[\frac{\rho \tau_t}{\rho-1} + \frac{(\nu-\rho)\tau_t}{(\nu-1)(1-\rho)} \right]$.

the pre-regulation equilibrium (see Appendix B.4 for a proof). As shown in Appendices B.4 and B.8, the LA estimator yields unbiased estimates of the average treatment effects on the regulated and unregulated firms' outcomes, as well as on aggregate emissions, as long as (1) there is no entry and exit of firms and (2) individual market shares remain approximately at their pre-regulation period values. However, if market shares adjust over time, the LA yields biased estimates of treatment effects, and the bias can be up or down.

3.3 Monte Carlo Experiments

We now perform a series of Monte Carlo experiments in order to demonstrate three important features of our GMM estimator relative to the alternative methods. First, our GMM estimator performs well in finite samples. Second, alternative estimators produce non-negligible biases. Third, our estimator still perform well even when the exclusion restrictions do not hold exactly.

To evaluate the finite sample properties of our GMM estimator and the two main alternatives, we generate simulated data sets as follows. A fixed number of firms could potentially operate in each of two periods – a pre-regulation period and a post-regulation period. Sector-wide expenditures are drawn randomly each period, as well as firm-level entry and exit decisions. In the first period, firms draw their productivity, number of plants, and α_{jfst} terms randomly. These variables completely determine unit cost, and hence prices and revenues, given the monopolistic competition assumption. We solve for plant-level emissions using the first order conditions (13). In the second period, we set $R_{jfst} = 1$ for all plants with pre-period emissions above a given threshold. Productivity updates according to a first-order Markov process. Revenues, market shares and plant-level emissions are solved again. We simulate the model for 24 different parameter combinations, varying $\rho \in \{.8, .9, .95\}$, $\mu^e \in \{0, .3, .5, .8\}$, and $\varrho \in \{0, .1\}$.³² For each parameter combination, we simulate 100 replications, with 10 sectors, 50 industries, and 500 firms per replication.

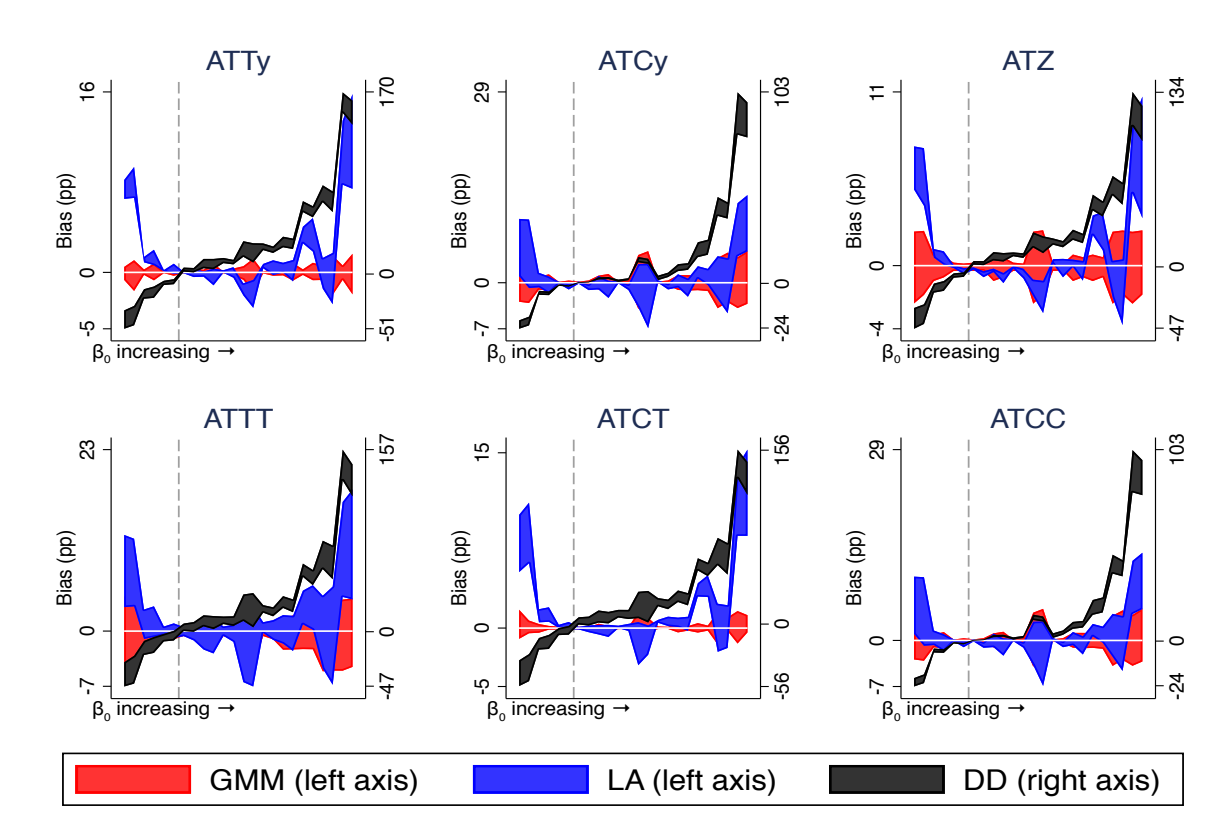
We then implement our estimation procedure, as well as multi-plant versions of the DD and LA estimators, on the simulated data.³³ For each estimation strategy, we compute firm-level ATT and ATC, plant-level ATTT, ATCT, ATCC, and aggregate effects on emissions for each simulated data set. Since we control the data generating process, we

³²In all combinations, we set $\nu = .3$, $\gamma = .8$, $\mu^z = .2$, and $\sigma = -0.1$.

³³To compute plant-level and aggregate effects on emissions for the LA method, we combine regression coefficients from the firm-level revenues regression (29) with the plant-level emissions regression (25).

can also compute the true treatment effect of the regulation on any given firm or plant, for any set of exogenous parameters. For each replication and estimator, we compute the difference between the estimated treatment effect and the true treatment effect and plot the interquartile range of this distribution in Figure 2.

Figure 2: Bias in Average Treatment Effects in Monte Carlo Experiments



Notes: Subfigures plot the interquartile range of the distribution of estimation errors by estimator and parameter combination for average and aggregate treatment effects across 100 replications. On the x-axis, parameter combinations are ordered by $\beta_0 = \frac{\rho\tau_t}{\rho-1}$. GMM and LA estimation biases are measured on the left axis, while DD estimation biases are measured on the right axis. In all subfigures, parameter combinations to the left (right) of the vertical dashed line indicate simulations for which $\tau_t > 0$ (< 0).

Our GMM procedure does recover unbiased estimates of treatment effects in our simulations. Figure 2 shows that our GMM procedure (plotted in red, measured on the left axis) yields small errors that are symmetrically centered on zero. This applies for all average and aggregate effects and across all parameter combinations.

By contrast, in Figure 2, we find that the DD and LA estimators produce substantial biases that can go in either direction, and even get the sign of the effect wrong. The DD

estimator (plotted in black, measured on the right axis) substantially underestimates the treatment effects on revenues when $\mu_t^e < \mu_t^z$, and overestimates it otherwise. Indeed, for most parameter combinations, the interquartile range of the errors does not contain zero at all. The errors for the LA estimator (plotted in blue, measured on the left axis) are smaller in magnitude. For some parameter combinations, the interquartile range covers zero, and even appears nearly centered. Indeed, when market shares move very little between the pre-regulation and post-regulation periods, a local approximation is appropriate. But in general, the LA method can yield substantially biased estimates, and the bias could go in either direction.

Finally, to provide an adversarial test of our GMM estimator, we generated comparable simulated data sets assuming firms engage in Bertrand-Nash pricing, as in Atkeson & Burstein (2008). Our GMM estimator, by contrast, relies on the assumption of monopolistic competition. The results, summarized in Appendix C.2, show that our GMM procedure continues to exhibit superior performance even in this setting, compared to DD and LA methods.

4 Application to the EU Emissions Trading System

In this section, we demonstrate our method by estimating the effect of the EU ETS on the revenues and emissions of French manufacturing firms.

4.1 Background and Data

The EU ETS is the European Union’s flagship climate policy. An EU-wide carbon market was first proposed in 2000, passed into law in 2003, and launched in 2005. The program would cap the combined carbon emissions of over 10,000 large power and manufacturing plants across Europe, while allowing inter-plant trading of emissions permits to keep compliance costs low. In manufacturing, the program included all combustion installations with a rated thermal input greater than 20MW, and other productive processes with capacity or output greater than predetermined industry-specific thresholds.

The EU ETS was implemented in three distinct trading phases, which differ somewhat in ambition and rules. Phase I, from 2005 to 2007, was designed as a trial period—the emissions cap was more generous, permits were allocated on the basis of historical emissions, but those permits could not be banked for future compliance. Phase II, from 2008 to 2012, ran concurrently with the first commitment period under the Kyoto Protocol, and the

emissions cap was set to meet the EU’s collective emission reductions commitment. Phase III, from 2012 to 2020, centralized the permit allocation process from national regulators to the European Commission, and increasingly allocated permits by auction.

Starting in Phase III, the European Commission created an integrated registry to keep track of the whole market. We retrieved the address of each regulated installation in France from this registry, as well as the initial date of regulation and the unique French tax identifier of the firm that owns the installation. We then matched EU ETS installations to plants using tax identifiers and street addresses (see Appendix D.3 for details). In total, we count 1,415 installations in France ever regulated under the EU ETS across 1,264 plants and 846 firms (note, a large plant may include multiple installations).

To build proxies for the CES price indices in our model and to estimate treatment effects, we need data on the revenues of the universe of firms. For this purpose, we use data reported to the French tax authority.³⁴ We classify firms into “industries” using a 4-digit activity code declared by the firms, and use the first 2 digits of the activity codes to define “sectors,” following Harrigan et al. (2018a). We report descriptive statistics for sales, employment, and number of plants for the year 2004 – the last pre-regulation year – in panel A of Table 1, in columns (1) and (2). Panel B reports emissions and employment at the plant level. The number of employees for each plant comes from the *Stock of Establishments*, while plant-level emissions are based on detailed fuel-consumption surveys (EACEI). The fuel-consumption data are collected annually for plants with more than 250 employees, while smaller plants are sampled randomly. More details about all of the data sets and processing can be found in Appendix D.

4.2 Joining Model With Data

Because our model necessarily portrays a somewhat simplified reality, there are three practical issues that must be addressed in joining the model with the data: (i) treatment endogeneity, (ii) foreign markets, and (iii) unobserved emissions for small plants.

Treatment endogeneity. From columns (1) and (2) of Table 1, we can see that the policy is incomplete – only 0.1% of French manufacturing firms operate any EU ETS regulated plants. We can also see that the policy targeted large firms – regulated firms

³⁴France’s statistical agency INSEE record these data in the FICUS database (for years 1994 - 2007) and the FARE database (2008 - 2016). Firms are identified with a unique tax identifier across both data sets. These data are confidential and access is subject to authorization by the *Comité du Secret Statistique*.

Table 1: Descriptive Statistics in 2004

	Full Sample			Matched Sample		
	regulated (1)	unregulated (2)	p-val (3)	regulated (4)	unregulated (5)	p-val (6)
<i>Panel A: Firm Level Dataset</i>						
Sales (Millions euros)						
Domestic	251.4	4.053	0.000	109.6	103.0	0.718
Export	98.65	1.033	0.000	68.50	52.76	0.172
Total	350.1	5.086	0.000	178.1	155.8	0.370
Market Share	0.066	0.001	0.000	0.033	0.029	0.336
# Workers	846.8	22.77	0.000	568.0	502.1	0.455
# Plants	3.625	1.160	0.000	2.353	2.624	0.279
α_{2004}	0.513	0.050	0.000	0.376	0.404	0.573
Δy	0.210	0.114	0.002	0.203	0.175	0.437
# Firms	363	149188		255	255	
<i>Panel B: Plant Level Dataset</i>						
CO ₂ Emissions ('000 Kg)	76.92	4.722	0.000	58.52	76.62	0.748
# Workers	403.4	172.3	0.000	392.1	422.7	0.702
# Firms	301	6657		156	157	
# Plants	606	7918		173	173	

Notes: Values indicate mean of annual observations for 2004 – the last pre-regulation year – by regulation status. Columns 1 and 2 includes all firms in all industries with 2-digit codes between 15 - 37 in the NAFRev.1, while columns 4 and 5 include only firms from the matched sample. Columns 3 and 6 report p-values from t-tests of difference in means between regulated and unregulated firms.

operate more plants, employ more workers, and have higher sales. If productivity followed a random walk around this period, then there would be no reason to expect changes in revenues to correlate with 2004 levels, and the association between pre-regulation size and treatment would not be a cause for concern. More generally, however, one might be worried that the outcomes for regulated and unregulated firms would have followed different trends even without the policy, given the clear correlation with pre-regulation size.

To address this concern, we match regulated firms to unregulated firms within each industry based on pre-regulation characteristics. Following Cabel & Dechezlepretre (2016), most studies of the EU ETS match on revenues within industries, and assume variation in the dispersion of economic activity across firms' plants generates exogenous variation in regulation status. In particular, it is usually assumed that firms with fewer, but larger, plants are more likely to be regulated under the EU ETS, and that this variation does not independently influence revenues. However, our model indicates precisely the opposite: unless $\sigma = 0$, the dispersion of economic activity across firms' plants directly influences

revenues. Hence, we would expect that two firms with the same pre-regulation revenues but different pre-regulation dispersion measures would grow at different rates post regulation, even if there were no treatment effect. For this reason, we match firms both on pre-regulation revenues and the pre-regulation dispersion measure $\alpha_{fis,2004}$. We also match on 2004 workers, 2004 domestic market share, and log export sales. While the model does not indicate that we need to match on these variables, in reality, one could imagine that firms with more workers (similarly, market share, export sales) grow at different rates compared to firms with less workers (market share, export sales), so we prefer to match on these variables as well to control as flexibly as possible for differential trends.³⁵

The descriptive statistics for the matched sample are reported in columns (4) and (5) of Table 1. We match 255 out of 363 regulated firms, each to a single unregulated firm. Because the EU ETS regulates many firms in the upper tail of the distribution, it is sometimes impossible to find closely matched unregulated firms within the same industries. The ones we match are therefore substantially smaller than the typical EU ETS firm, with an average 2004 market share of 3% instead of 6%. The matched regulated and unregulated firms are balanced in terms of the key variables: revenues, revenue growth, and dispersion. They are less well-balanced in terms of exports, but we show later that our results are robust to re-matching on exports (see Section 4.5).

Given that EU ETS regulation is based on plant size, it might seem that matching on both firm-level revenues and economic dispersion across plants would eliminate all identifying variation. The critical insight is that the model requires only that we condition on a *particular* measure of economic dispersion. The identification strategy remains valid as long as there exists some alternative measure of dispersion that predicts treatment, conditional on $\alpha_{fis,2004}$. Appendix Table A.1 shows that, in fact, across-plant concentration as measured by the Herfindahl index, $\sum_{j \in \Omega_{fis,2004}} \alpha_{jfis,2004}^2$, correlates with future regulation status, even conditional on $\alpha_{fis,2004}$.³⁶ Thus, even after matching on both revenues and economic dispersion, there remains variation in treatment that can be used to identify the effect of the policy.

To estimate μ^Z , we estimate equation (25) by OLS. Having taken Δy_{fist} to the left hand

³⁵In practice, we use the coarsened exact matching program from Blackwell et al. (2009), matching exactly on industry, and within strata on revenues in 2004, $\alpha_{fis,2004}$, market share in 2004, total number of workers in 2004, and average revenue growth prior to 2004.

³⁶An obvious alternative measure of concentration, in this context, would be to measure the difference between each plant's capacity and the activity-specific regulation thresholds – this would directly capture that aspect of concentration that is the basis of treatment. Unfortunately, data on these plant-specific capacities are not available.

side, the model indicates that no further matching is necessary to address endogeneity of the regulation at the plant level. Hence, restricting the sample to matched regulated and unregulated plants is not necessary. Even so, we evaluate the robustness of our findings to using a sample of plants matched on their pre-period CO₂ emission levels and growth rates. As shown in Panel B of Table 1, regulated plants have more employees and higher emissions. Yet, because the regulation is based on production capacity rather than emissions, some high-emitting plants are not regulated. Figure D.10 plots the average regulation status by 2004 CO₂ emissions levels, and shows there is no clear threshold above which regulation is complete. This is what allows us to match at the plant level.

Foreign markets. Our empirical strategy relies on constructing empirical counterparts to the theoretical CES price index, which requires information on all competing firms. While the model references a single market, in reality French firms compete in an international economy. Firms export some of their output into foreign markets, and they compete with imports from those foreign markets. We need some way of accounting for these exports and imports in our price indices, without radically expanding the data requirements to include all firms in the world.

To account for exports, we distinguish between domestic and foreign sales. For most of the analyses below, we take domestic sales of French firms as the outcome variable. Domestic sales represents about 75% of total French-firm revenues, depending on the year of the sample (see Figure D.13, left). We can only compute counterfactual revenues using our method for domestic sales, but we also sometimes present estimates on total sales for comparison with other methods. We also mostly restrict our analysis to CO₂ emissions generated in production for domestic sales, which we compute by multiplying total emissions by the domestic revenue share. Figure D.13 (right) shows that our measures of both total and domestic emissions track trends in the National Emissions Inventory reasonably well.

Imports present a greater challenge. Imports account for roughly one third of French consumption of manufacturing goods (see Figure D.11), but we do not observe the market shares of individual foreign sellers in the French market – information needed to build the CES price indices. To sidestep this missing data problem, we use the bilateral trade flows recorded in the BACI dataset (Gaulier & Zignago, 2010), and assume that each origin-country-industry is associated with only one firm exporting to France. We then investigate the sensitivity of our findings to the alternative assumption that we can leave

foreign imports out of the computation of treatment effects entirely. This alternative would be valid if there is separability between the CES price index for domestic varieties and the one for foreign varieties.

Imputing emissions for small plants. One quantity of interest is how the regulation affected aggregate CO₂ emissions from French manufacturing. However, the fuel-consumption EACEI survey only includes 10% of all plants each year, omitting many unregulated plants especially (see Figure D.9). To estimate the aggregate effect based on this sample, we need two further assumptions.

First, we assume there are no systematic differences between surveyed and unsurveyed plants, conditional on observables. This implies that the treatment effects estimated on the sample are valid for the whole population of plants.

Second, we impute the emissions for unsurveyed plants by multiplying the plant-level revenue by the median sector-year emissions per unit of revenue. This imputation relies on taking literally the modelling assumption that the emission intensity of revenues does not vary across plants within a sector-year, absent the regulation.³⁷ See Appendix D.6 for details on this imputation. We will report results for the EU ETS's effect on total emissions based on surveyed data only and on these imputed data. In the latter case, we remain alert to the fact that the estimated aggregate effect on emissions relies on stronger assumptions than the rest.

4.3 Evaluating the Underlying Assumptions of our Model

Our GMM procedure generalizes prior approaches in at least three important ways – it accounts for firms' multi-plant structure, it accounts for spillovers due to imperfect competition, and it accounts for market shares varying over time (including from entry and exit). Before reporting on our results, we first discuss why each of these generalizations is particularly valuable for studying a policy like the EU ETS.

Multi-plant firms. We have explicitly modelled the multi-plant structure of firms and argued that this is critical for causal identification. If firms optimize production over their entire set of plants, representing regulation as a simple binary variable will give rise to omitted variable bias. In particular, our estimating equation for firm-level outcomes ought

³⁷In practice, the observed emission intensity of revenues does vary, but this could be due to measurement error.

to include a continuous treatment intensity variable, as well as a measure of dispersion of activities across plants (see equation 22). It is easy to verify that this multi-plant model provides additional explanatory power over a model with only a binary treatment variable.³⁸

We can further assess the importance of the multi-plant structure as well as matching by estimating the multi-plant DD model (equation 27). Figure 3 reports estimates of β_0 and β_1 resulting from estimating this model via OLS. We also add firm-level controls for pre-treatment revenues, revenue growth, and dispersion. The model is estimated in long differences relative to the base year 2004.³⁹ In the left panel of Figure 3 (using the full sample), we find that β_0 is positive and statistically distinguishable from zero in the pre-treatment periods, and negative in the post-regulation period (though not statistically significant). Firms that had grown faster, and were therefore larger by 2005, were more likely to be regulated. If productivity follows a general autoregressive process, past productivity would also predict future revenues. The negative β_0 s in the post-treatment period could therefore be the result of a negative effect of regulation, or of larger firms tending to grow slower.

The right panel of Figure 3, which reports on the matched sample, helps to resolve this ambiguity. We verify that β_0 is not significantly different from zero in the pre-period and find that it tends to be positive in the post-treatment period. Having addressed selection on productivity through matching, these unbiased post-treatment estimates of β_0 indicate a positive direct effect of the regulation. Therefore, the negative β_0 s in the full sample are best explained by productivity not following a random walk. This demonstrates the necessity of matching.

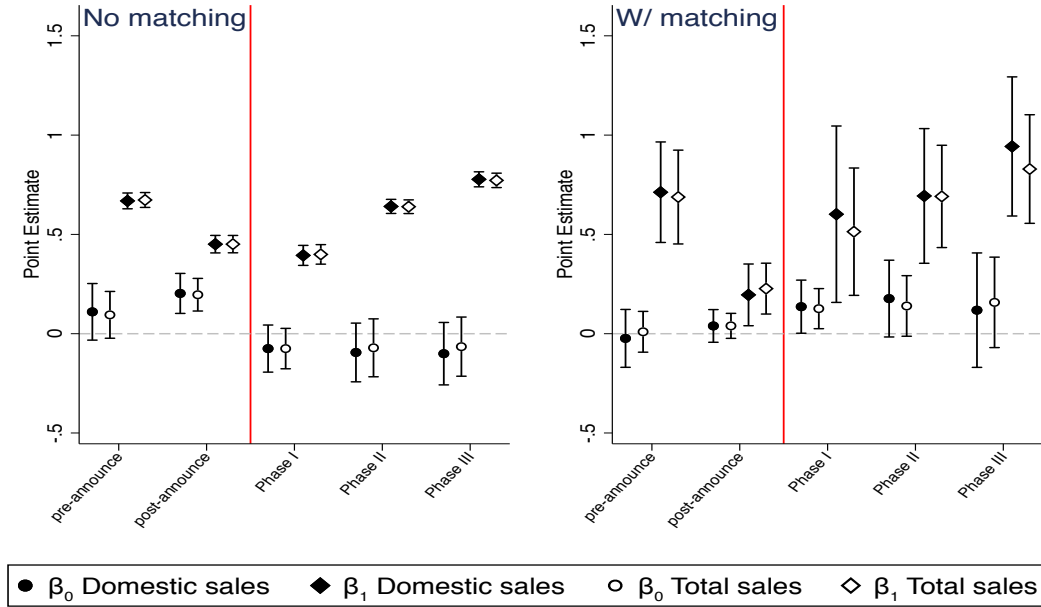
The coefficients on economic dispersion, β_1 , are non-zero, both in the full and matched samples. Even conditional on treatment, firms with more dispersed production tended to grow faster. Since the measure of dispersion is correlated with regulation by construction – the labor shares appear in both terms – failing to control for it will result in omitted variable bias.

The multi-plant structure is also important for studying plant-level outcomes, such as

³⁸The multi-plant model is associated with lower values for the AIC and BIC, and higher values for R^2 , for all three phases of the EU ETS.

³⁹We pool years to estimate β_0 and β_1 once for each period—the pre-announcement period (1994-1999), the post-announcement period, which is still prior to implementation (2000-2003), and each of the three trading phases. For ease of interpretation, we invert the outcome variable for years prior to 2004, so that the dependent variable represents the growth rate between a pre-period year and 2004. We similarly invert the ratio for $\Delta\alpha_{*fst*}$ for pre-2004 years. All regressions include industry-by-year fixed effects, and standard errors are clustered on the industry.

Figure 3: Firm-level DD Estimates for Revenues




Notes: Markers indicate estimates of β_0 and β_1 from equation (27) using domestic sales (in black) or total sales (in white) as the dependent variable. Bars indicate 95% confidence intervals. Left (right) panel includes all (matched) firms. For each period, we estimate a pooled regression, combining all years included in the period. The vertical line indicates the beginning of the regulation period. All regressions include industry-by-year fixed effects. Standard errors are clustered on the industry.

emissions. When we estimate a simple DD model for emissions at the plant-level (equation 28), it appears as though the regulation not only affected regulated plants (Figure 4, left panel), but also unregulated plants operated by regulated firms (Figure 4, right panel). Estimating the response only for regulated plants, and ignoring the spillovers on other plants owned by the same firms, will not be able to identify the true environmental effect of the policy.⁴⁰

Spillovers. Another critical feature of our model is the presence of spillovers from regulated to unregulated firms. Spillovers will be negligible in some applications, such as when the outcome variable is not strongly affected by competition, or when the treated firms account for a small share of the market. These conditions clearly do not hold when studying the EU ETS’ effects on revenues and emissions. Our model – indeed, any model of imperfect competition – shows that firms’ revenues and emissions are inter-dependent.

⁴⁰These estimates reveal similar “reallocation effects” as estimated by Gibson (2019) and Soliman (2020), showing counterbalancing changes in emissions across regulated and unregulated plants owned by regulated firms in the US context of the Clean Air Act amendments.

Figure 4: Plant-level DD Estimates for Emissions



Notes: Left (right) panel presents estimates of $\beta_0^{Z,DD,TT}$ ($\beta_0^{Z,DD,CT}$) derived from OLS estimation of equation (28). Bars indicate 95% confidence intervals. For each period, we estimate a pooled regression, combining all years included in the period. The vertical line indicates the beginning of the regulation period. All regressions include industry-by-year fixed effects. Standard errors are clustered on the industry.


Moreover, EU ETS firms in our sample command, on average, over 6% of the market in their industry (in 2004). In this setting, the independence condition required for the DD estimator to identify the effects of an incomplete regulation (i.e. SUTVA) is violated.

Both our GMM estimator and the LA estimator are flexible enough to allow for spillovers within an industry and across industries. Figure 5 presents the GMM estimates (in red) and the LA estimates (in blue) for the matched sample.⁴¹ The estimated values of β_0 are consistently positive, indicating a positive direct effect of regulation. Meanwhile, β_2 tends to be negative, and more precisely estimated with GMM than the LA method, which indicates that firms grow slower when there is a higher share of regulated firms in

⁴¹As before, we also control for the vector of firm characteristics, \mathbf{X}_{fist_0} , which includes the pre-period level and growth rate of revenues and economic concentration at the firm-level. For the moment conditions at the industry-level, we control for a vector of industry pre-period characteristics, including total sales in level and growth rate and average concentration at t_0 . Standard errors are clustered to allow for arbitrary correlation within the industry and over time.

their industry. This is a tell-tale sign of within-industry spillovers.^{42,43}

Figure 5: Reduced-Form Coefficients β_0 , β_1 , β_2 , and β_z



Notes: Left panel presents the estimates of β_0 , β_1 , and β_2 from our GMM estimator (in red) and from the LA model (in blue). Right panel presents estimates of β_z derived from OLS estimation of equation (25). Bars indicate 95% confidence intervals. For each period, we estimate a pooled regression, combining all years included in the period. The vertical line indicates the beginning of the regulation period.

Market share stability. The LA method can be shown to be a local approximation of our GMM estimator around the pre-treatment equilibrium. This local approximation has one key advantage – it can be estimated without knowing market shares of all firms in the market (just the combined market share of regulated firms). Moreover, our Monte Carlo experiments show that the LA estimator performs well in situations where the market shares are fairly stable. The LA method could therefore be a valuable tool in data-constrained environments with a fairly stable set of firms, and where the incomplete regulation is expected to have a moderate treatment effect.

In our application, as in many developed country contexts, we observe revenues for the entire universe of French firms. We are also interested in the effects of a substantial policy experiment, over a long enough time period, that one may prefer not to assume *a priori* that market shares are stable. Indeed, Figure A.2 shows substantial differences between pre-treatment market shares $\theta_{fi,2004}$, and the theoretically-consistent Sato-Vartia weights

⁴² $\beta_2 \neq 0$ is sufficient, but not necessary, to indicate the existence of spillovers. Spillovers within industries could exactly counterbalance spillovers across industries, which would result in $\beta_2 = 0$, or equivalently, $\rho = \nu$. Across-industry spillovers may be small in our application, as indicated by the sum of β_0 and β_2 ($\beta_0 + \beta_2 = \frac{\nu\tau_t}{1-\nu}$) being close to zero.

⁴³Note, also, that both the GMM and LA methods return negative estimates of β_z , consistent with the assumption in our model that $\mu^z > 0$.

ϕ_{fit} , indicating that market shares indeed move substantially over time.⁴⁴

4.4 Results

Revenues. Figure 6 summarizes the estimated treatment effects for the revenues and emissions of French manufacturing firms using our GMM estimator (in red), alongside the results for the DD (in black) and LA (in blue) estimators for comparison. We estimate these effects in long differences relative to 2004, which is the last pre-regulation year. The underlying model parameters are estimated separately for each of the three EU ETS trading phases (see figure 5), then used to compute firm-level counterfactuals under the scenario that $R_{jfst} = 0$ for all plants, and finally aggregated into average treatment effects for each year.

The top-left panel of Figure 6 plots the EU ETS's average effects on the domestic sales of regulated firms. With the GMM estimator, we find that the EU ETS increased domestic sales for the average regulated firm by 6-9% annually, relative to the unregulated counterfactual. The DD estimator overstates this effect by 1-2 percentage points, and the LA estimator understates the effect by 2-5 percentage points, depending on the year. In the aggregate, our GMM estimates imply that the EU ETS increased domestic revenues for regulated firms between 1-3% as a group.

We also find a negative effect of the regulation on the domestic sales of the average unregulated firm (see top-right panel of Figure 6). Our GMM estimates imply a 0.2-0.7% reduction in domestic revenues for unregulated firms, as a group. By contrast, with the LA model, we estimate that regulation had almost no effect on the sales of unregulated firms through phase I, and then a modest *positive* effect in phases II and III. For the DD estimator, the effect on unregulated firms is fixed at zero by assumption.

Our results suggest that previous estimates of the EU ETS's effect on revenues have probably been off by a few percentage points, errors that may be meaningful to a social planner. However, we learn a lot more from our new estimates than this difference in magnitude. We have shown that the DD and LA estimators are subject to substantial and unpredictable biases, even to the extent they can get the incorrect sign of the treatment effects. Even if the GMM procedure had yielded identical results in the present application, then, it would not be possible to judge the credibility of prior estimates except by reference to our findings.

⁴⁴Looking directly at market shares, we find that fully three quarters of firm-year observations differ from 2004 by more than 10%, and half differ by more than 25%.

Emissions. The middle and lower panels in Figure 6 summarize the EU ETS's estimated effects on domestic emissions.⁴⁵ Our GMM estimates indicate that the EU ETS reduced emissions at the average regulated plant between 5 -25%, depending on the year (see middle-left panel). These findings suggest that the technique effect dominated in all three phases, resulting in a decrease in emissions. The LA estimator provides broadly similar results, but systematically overstates the emissions reductions by a few percentage points. The DD estimator, by contrast, generates smaller effects on emission reductions during phases I and II, and a modest *increase* in emissions in phase III.

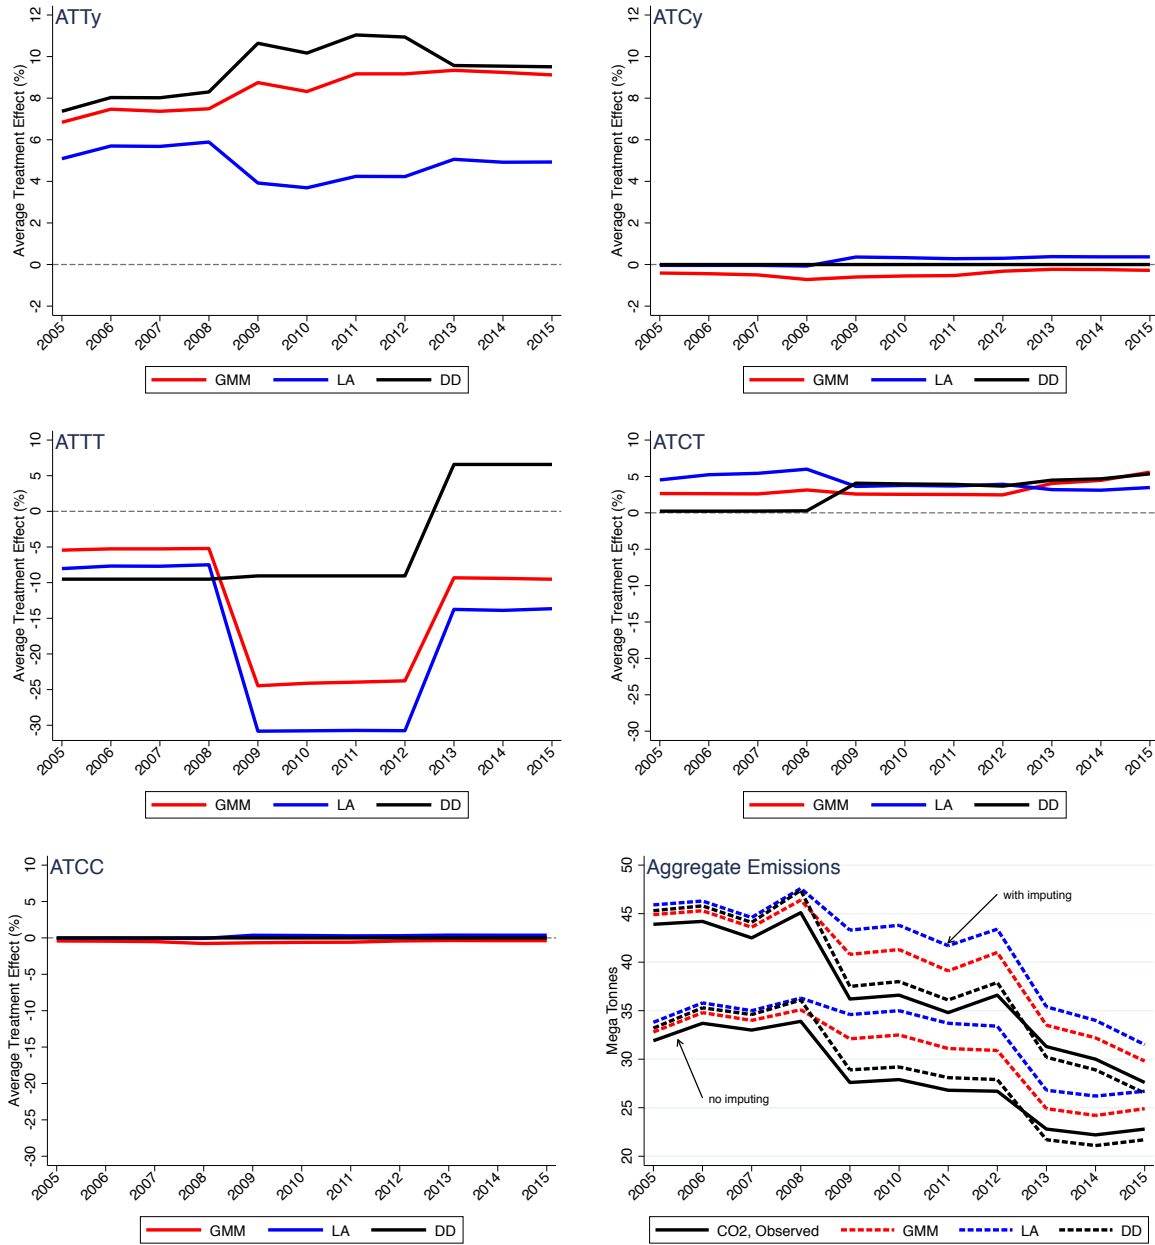
We also find an increase in emissions at the control plants operated by treated firms (middle-right panel). This suggests that regulation-induced cost reductions at regulated plants increased revenues of regulated firms, and thereby emissions at unregulated plants owned by regulated firms. The LA and DD estimators sometimes overstates these effects and sometimes understates them, relative to our GMM estimator. Furthermore, all three estimators agree that regulation had practically no effect on the emissions of plants owned by unregulated firms (bottom-left panel). Yet, we observe a slight negative effect using our GMM estimator.

Finally, the bottom-right panel of Figure 6 plots the aggregate observed CO₂ emissions generated in the production of goods for the French market, along with the counterfactual levels computed using the GMM, LA, and DD estimators. The bottom set of trajectories only includes plants observed in the fuel-consumption surveys, whereas the top set of trajectories includes all plants with imputed emissions for small unsurveyed plants.

Without imputing, we find that the regulation lowered domestic emissions between 0.9-4.6 million tonnes annually, or between 3-16% of total (non-imputed) emissions. For robustness, we compare our aggregate results to the ones obtained with imputing emissions for unsurveyed plants. With imputing, we find that the regulation lowered emissions between 1.0-4.6 million tonnes annually, depending on the year, or between 2-12% of total emissions. By comparison, we find that the LA estimator is biased upward, while the DD estimator is biased downward, the latter yielding virtually zero effect during phases I and II and a small increase in phase III. In aggregate, we find that the EU ETS reduced emissions generated in the production of goods for domestic consumption between 2005 and 2015 by 28.0 (29.1) million tonnes without (with) imputing.

⁴⁵Note that our results here refer to the portion of emissions associated with production for the domestic market.

Figure 6: Treatment Effects on Domestic Firm-Level Revenues, Plant-Level Emissions and Aggregate Emissions



Notes: Figure presents the average effects on revenues for treated T (top left) and control C (top right) firms, and on plant-level emissions for TT plants (middle left), CT plants (middle right) and CC plants (bottom left) generated for the domestic market (using plants in the EACEI surveys). Bottom right panel presents observed total CO₂ emissions generated for the domestic market (solid black) along with counterfactual emissions computed using GMM, LA and DD methods. Lower trajectories include only plants observed in the EACEI survey, while top trajectories impute emissions for small missing plants.

4.5 Robustness Checks

In this section, we explore the robustness of our results to the various data choices that were made to accommodate factors outside of our model.

First, lacking data on firm-level imports to France, we have assumed so far that the set of foreign firms shipping to France within an industry-country-of-origin is always singleton. We conduct a robustness test under the alternative assumption that we can leave foreign imports out of the computation of treatment effects entirely. This exercise amounts to assuming Cobb-Douglas aggregation between domestic consumption and foreign consumption at the sector level, so that the share of expenditures on domestic goods would be fixed. Under this assumption, industry-level CES price indices on the French market can be separated into a CES price index for domestic varieties and a CES price index for foreign varieties, and the two are completely independent. As a result, we could construct the CES price index for domestic varieties using only French firms and ignoring international imports. This is implicitly the approach taken by previous work based on local approximations (Cai & Szeidl, 2022; Muehlegger & Sweeney, 2021; Rotemberg, 2019). Appendix Figures A.3-A.4 show that results are robust to this alternative assumption.

Second, we explore the robustness of our results to alternative matching criteria. We prefer to match on as many covariates as possible, but the model indicates that only revenues and across plant dispersion need to be controlled for when choosing a matched sample. We repeat the exercise matching only on those variables that the model indicates should be match on. Not surprisingly, covariate balance deteriorates in this alternative matching procedure, but results remain broadly consistent with our preferred specification (Appendix Figure A.5).

Finally, we explore robustness to alternative definitions of competitors sets. In our baseline specification, we treat the 4-digit NAFRev.1 classification of the firm as the firms “industry”. It seems reasonable that firms that declare the same 4-digit industry code compete with each other, but it certainly could be possible that firms directly compete with firms that declare other codes as well. To explore this hypothesis, we define the industry alternatively as the 3-digit NAFRev.1 code. Sectors are defined as above. In Figure A.6, we find that treatment effects on emissions are broadly similar to the baseline specification, though effects on domestic revenues are smaller.

5 Conclusion

When firms compete imperfectly, changes in the production costs of one firm affect the output decisions of other firms. These across-firm dependencies are inconsistent with a necessary assumption in difference-in-differences estimation – namely, the stable unit treatment value assumption. We show that, if the data generating process coincides with standard modeling assumption from industrial organization and international trade, the DD estimator tends to be informative about the sign of the cost effect of regulation, but neither the sign nor the magnitude of the average effects on firm-level revenues and emissions. In our view, this result implies that DD should not be relied upon for inference with respect to average treatment effects of regulation – and more broadly, cost shocks – on firms’ revenues and emissions.

We build an estimation process that is consistent with standard modeling assumptions and allows researchers to recover average and aggregate treatment effects of regulation, relying purely on panel data information on market shares (and emissions, if so desired), and the vector of regulation. The procedure allows for endogenous technological adoption, the effect of which would be hard to predict *ex ante*. The procedure generalizes previous work based on local approximations around the pre-regulation equilibrium. Hence, our procedure can be applied even if market shares move substantially over the study period. We also show how to account for multi-plant production, non-random regulation, and general productivity growth processes.

We find that, contrary to fears of being put at a competitive disadvantage, the flagship EU climate policy did not increase the costs of regulated French manufacturers. Rather, we find that revenues of regulated French manufacturers increased between 6% and 9% as a result of the EU regulation, and that emissions at regulated plants fell between 5% and 25%, depending on the year. Unregulated plants owned by regulated firms increased emissions slightly, but not enough to counteract the emissions reductions at regulated plants. Aggregate emissions generated in France for the French market fell between 3% and 16% relative to the unregulated counterfactual.

Though regulated firms increased market share, this does not necessarily mean that regulated firms benefited from the policy. Firm-level benefits are measured in profits, not sales. We do not know the investment costs incurred to lower costs. A revealed preference argument would lead one to believe that, if firms did not undertake these cost-reducing investments *ex ante* to the regulation, then they were not expected to be profitable. Hence, it is possible that total profits of regulated firms fell due to investment costs. Conversely,

it is also possible that x-inefficiencies or other forms of inertia kept managers from making profitable investments until the regulation took effect. Without detailed data on investment and a dynamic model, we cannot discriminate between these two explanations and we leave this investigation for further research.

The results are consistent with a Porter effect, in which regulation induces technological improvements that lower both emission intensity and costs. This outcome may have been triggered by specific measures taken at the European and national levels to avoid hurting domestic firms' competitiveness and to favor technological adoption. This does not mean that the EU ETS will continue to produce both revenue growth and emissions reductions in the future, of course, nor that environmental regulations more generally can be expected to trigger Porter effects. In general, the environmental benefits of regulation should be weighed against the economic costs, which could include erosion of domestic industry and distributional effects on consumer welfare.

References

- Acemoglu, D. & Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. *Handbook of Labor Economics*, 4(B), 1043–1171.
- Adao, R., Arkolakis, C., & Esposito, F. (2020). General Equilibrium Effects in Space: Theory and Measurement. *National Bureau of Economic Research Working Paper*, 25544.
- Alfaro, L., Conconi, P., Fadinger, H., & Newman, A. F. (2016). Do Prices Determine Vertical Integration? *Review of Economic Studies*, 83(3), 855–888.
- Ambec, S., Cohen, M. A., Elgie, S., & Lanoie, P. (2013). The Porter Hypothesis at 20: Can Environmental Regulation Enhance Innovation and Competitiveness? *Review of Environmental Economics and Policy*, 7(1), 2–22.
- Atalay, E., Hortacsu, A., & Syverson, C. (2014). Vertical Integration and Input Flows. *American Economic Review*, 104(4), 1120–48.
- Atkeson, A. & Burstein, A. (2008). Pricing-to-Market, Trade Costs, and International Relative Prices. *American Economic Review*, 98(5), 1998–2031.
- Baier, S. L. & Bergstrand, J. H. (2009). Bonus vetus OLS: A simple method for approximating international trade-cost effects using the gravity equation. *Journal of International Economics*, 77(1), 77–85.
- Barrows, G. & Ollivier, H. (2018). Cleaner firms or cleaner products? How product mix shapes emission intensity from manufacturing. *Journal of Environmental Economics and Management*, 88(134–158).

- Bartram, S. M., Hou, K., & Kim, S. (2022). Real Effects of Climate Policy: Financial Constraints and Spillovers. *Journal of Financial Economics*, 143(2), 668–696.
- Bergquist, L., Faber, B., Fally, T., Hoelzlein, M., Miguel, E., & Rodriguez-Clare, A. (2022). Scaling Agricultural Policy Interventions: Theory and Evidence from Uganda. *Mimeo*.
- Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, 63(4), 841–890.
- Blackwell, M., Iacus, S., King, G., & Porro, G. (2009). cem: Coarsened exact matching in stata. *The Stata Journal*, 9(4), 524–546.
- Borusyak, K., Dix-Carneiro, R., & Kovak, B. K. (2022). Understanding migration responses to local shocks. *Mimeo*.
- Cai, J. & Szeidl, A. (2022). Indirect Effects of Access to Finance. *National Bureau of Economic Research Working Paper*, 29813.
- Calel, R. (2020). Adopt or Innovate: Understanding Technological Responses to Cap-and-Trade. *American Economic Journal: Economic Policy*, 12(3), 170–201.
- Calel, R. & Dechezlepretre, A. (2016). Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market. *Review of Economics and Statistics*, 98(1), 173–191.
- Calligaris, S., D’Arcangelo, F. M., & Pavan, G. (2022). The Impact of the European Carbon Market on Firm Productivity: Evidence from Italian Manufacturing Firms. *FEEM Working Paper*, 24.
- Colmer, J., Martin, R., Muûls, M., & Wagner, U. J. (2021). Does Pricing Carbon Mitigate Climate Change? Firm-Level Evidence from European Union Emissions Trading Scheme. *Mimeo*.
- Costinot, A., Donaldson, D., & Smith, C. (2016). Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from 1.7 Million Fields around the World. *Journal of Political Economy*, 124(1), 205–248.
- de Chaisemartin, C. & D’Haultfoeuille, X. (2022). Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey. *NBER Working Paper*, 29691.
- De Loecker, J. & Syverson, C. (2021). An industrial organization perspective on productivity. In *Handbook of Industrial Organization*, volume 4 (pp. 141–223). Elsevier.
- Dechezleprêtre, A., Nachtigall, D., & Venmans, F. (2018). The joint impact of the European Union emissions trading system on carbon emissions and economic performance. *OECD Economics Department Working Papers*, 1515.

- Fabra, N. & Reguant, M. (2014). Pass-Through of Emissions Costs in Electricity Markets. *American Economic Review*, 104(9), 2872–2899.
- Feenstra, R. C. (1994). New Product Varieties and the Measurement of International Prices. *American Economic Review*, 84(1), 157–177.
- Forslid, R., Okubo, T., & Ulltveit-Moe, K. H. (2018). Why are firms that export cleaner? International trade, abatement and environmental emissions. *Journal of Environmental Economics and Management*, 91, 166–183.
- Fowlie, M., Reguant, M., & Ryan, S. P. (2016). Market-Based Emissions Regulation and Industry Dynamics. *Journal of Political Economy*, 124(1), 249–302.
- Franklin, S., Imbert, C., Abebe, G., & Mejia-Mantilla, C. (2021). Urban Public Works in Spatial Equilibrium: Experimental Evidence from Ethiopia. *Mimeo*.
- Ganapati, S., Shapiro, J. S., & Walker, R. (2020). Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes. *American Economic Journal: Applied Economics*, 12(2), 303–42.
- Gaulier, G. & Zignago, S. (2010). BACI: International Trade Database at the Product-Level. The 1994-2007 Version. *CEPII research center Working Papers*, 23.
- Gibson, M. (2019). Regulation-Induced Pollution Substitution. *The Review of Economics and Statistics*, 101(5), 827–840.
- Greenstone, M. (2002). The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the Census of Manufactures. *Journal of Political Economy*, 110(6), 1175–1219.
- Harrigan, J., Reshef, A., & Toubal, F. (2018a). *Techies, trade, and skill-biased productivity*. Technical report, National Bureau of Economic Research.
- Harrigan, J., Reshef, A., & Toubal, F. (2018b). Techies, Trade and Skill-Biased Productivity. *National Bureau of Economic Research Working Paper*.
- Head, K. & Mayer, T. (2022). Poor Substitutes? Counterfactual methods in IO and Trade compared. *CEPR Discussion Paper*, DP16762.
- Helpman, E., Melitz, M., & Rubinstein, Y. (2008). Estimating trade flows: trading partners and trading volumes. *Quarterly Journal of Economics*, 123(2), 441–487.
- Hintermann, B. (2017). Market Power in Emission Permit Markets: Theory and Evidence from the EU ETS. *Environmental and Resource Economics*, 66(1), 89–112.
- Holladay, S. J. (2016). Exporters and the environment. *Canadian Journal of Economics/Revue canadienne d'économique*, 49(1), 147–172.

- Jaraite, 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).
- Joltreau, E. & Sommerfeld, K. (2018). Why Does Emissions Trading Under the EU ETS not Affect Firms' Competitiveness? Empirical Findings from the Literature. *Climate Policy*, 19(4), 453–471.
- Löschel, A., Lutz, B. J., & Managi, S. (2019). The Impacts of the EU ETS on Efficiency – An Empirical Analyses for German Manufacturing Firms. *Resource and Energy Economics*, 56, 71–95.
- Marin, G. & Vona, F. (2021). The Impact of Energy Prices on Socioeconomic and Environmental Performance: Evidence from French Manufacturing Establishments, 1997-2015. *European Economic Review*, 135(103739).
- Martin, L. A., Nataraj, S., & Harrison, A. E. (2017). In with the big, out with the small: Removing small-scale reservations in India. *American Economic Review*, 107(2), 354–86.
- Martin, R., Muûls, M., Preux, L. B. D., & 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.
- Martin, R., Muûls, M., & Wagner, U. J. (2015). The impact of the European Union Emissions Trading Scheme on regulated firms: What is the evidence after ten years? *Review of Environmental Economics and Policy*, 0(0), 1–21.
- Muehlegger, E. & Sweeney, R. L. (2021). Pass-Through of Own and Rival Cost Shocks: Evidence from the US Fracking Boom. *Review of Economics and Statistics*, 104(6), 1361–1369.
- Neumark, D. & Simpson, H. (2015). Place-based policies. In *Handbook of regional and urban economics*, volume 5 (pp. 1197–1287). Elsevier.
- 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.
- Rotemberg, M. (2019). Equilibrium effects of firm subsidies. *American Economic Review*, 109(10), 3475–3513.
- Sato, K. (1976). The ideal log-change index number. *The Review of Economics and Statistics*, (pp. 223–228).
- Shapiro, J. S. & Walker, R. (2018). Why is pollution from US manufacturing declining? The roles of environmental regulation, productivity, and trade. *American Economic Review*, 108(12), 3814–54.
- Soliman, F. (2020). Intrafirm leakage. *CESifo Working Paper*, 8619.

Vartia, Y. O. (1976). Ideal log-change index numbers. *Scandinavian Journal of Statistics*, (pp. 121–126).

For Online Publication

Estimating the Effects of Industrial Regulation when Treated and Control Firms Compete: A New Method with Application to the EU ETS

by Geoffrey Barrows, Raphael Calel, Martin Jégard, Hélène Ollivier


A Appendix: Further Empirical Results

Table A.1: Predicting Regulation Status

	Probit		OLS			
	everR (1)	everR (2)	everR (3)	everR (4)	maxR (5)	maxR (6)
$\ln y_{2004}$	0.614*** (0.030)	0.610*** (0.030)	.0041*** (0.00009)	.0039*** (0.00009)	.0035*** (0.00008)	.0024*** (0.00008)
$\Delta y_{pre2004}$	0.220*** (0.075)	0.222*** (0.074)	.0035*** (0.00024)	.0034*** (0.00024)	.0029*** (0.00021)	.0028*** (0.00021)
$\alpha_{f,2004}$		0.282 (0.075)		0.0412*** (0.248)		0.0154*** (0.0019)
$\sum_{j \in \Omega_{f,2004}} \alpha_{jf,2004}^2$	0.336** (0.155)	0.899* (0.521)	-0.0167*** (0.0011)	0.0554*** (0.0039)	-0.0042*** (0.0009)	0.0227*** (0.0035)
# obs	25,871	25,871	139,950	139,950	139,950	139,950


Notes: Table reports probit (columns 1 and 2), and OLS (columns 3-6) estimation of correlations between firm-level characteristics in 2004 and future regulation. Outcome variable in columns 1-4 is indicator variable for whether a firm ever operates a regulated plant. Outcome variable in columns 5-6 is the maximum value of $R_{fist} \in [0, 1]$ observed for the firm. All regressions include industry fixed effects. Columns 1-2 include only industries with some ever-regulated firms. Columns 3-6 include all industries. Standard errors are clustered at the industry level.

Figure A.1: Entry and Exit




Notes: Markers indicate point estimates resulting from OLS estimation of the effect of base-year-weighted industry treatment density on the log difference in the number of firms entering (white) or exiting (black) an industry in year t relative to 2004. Bars indicate 95% confidence intervals. We estimate long difference regressions year by year, with sector-by-year fixed effects. Standard errors are clustered on the sector.

Figure A.2: Market Share Stability



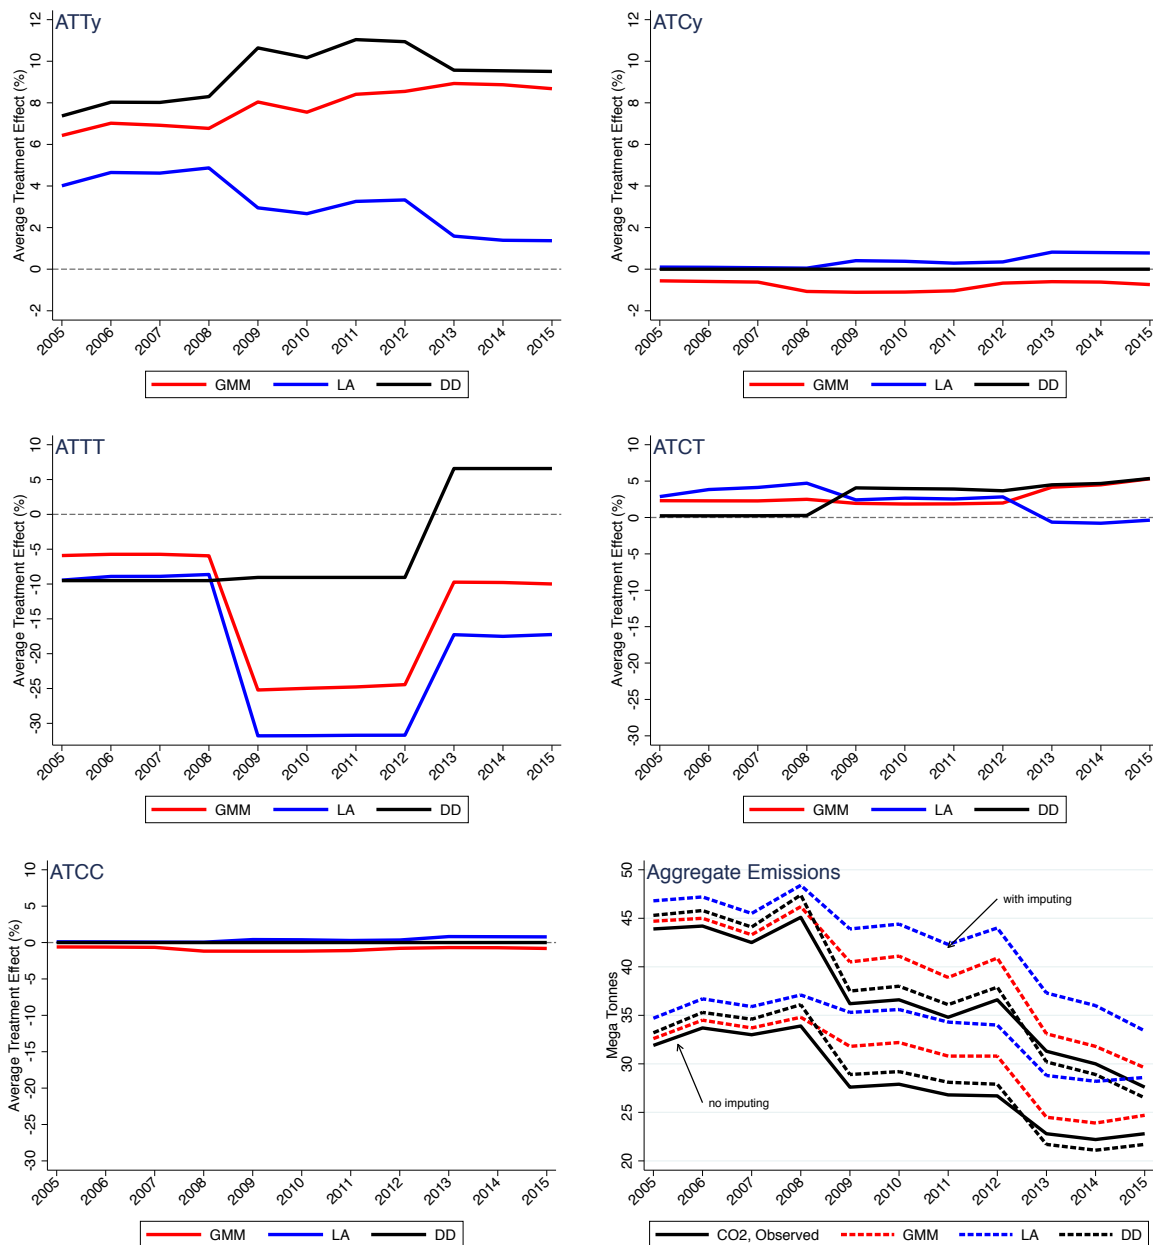
Notes: Left Figure plots Sato-Vartia weight ϕ_{fit} against market share in 2004, $\theta_{fi,2004}$. Right Figure plots the cumulative distribution of the ratio $\phi_{fit}/\theta_{fi,2004}$. The right tail of the figure has been truncated for ease of viewing.

Figure A.3: Point Estimates, LA Excluding Foreign Imports




Notes: Markers indicate point estimates resulting from OLS estimation of equation 29, excluding foreign imports, taking total sales (red) and domestic sales (blue) as outcomes. Bars indicate 95% confidence intervals. For each period, we estimate a pooled regression, combining all years included in the period. All regressions include sector-by-year fixed effects. Standard errors are clustered on the industry.

Figure A.4: Robustness Check, Excluding Foreign Imports



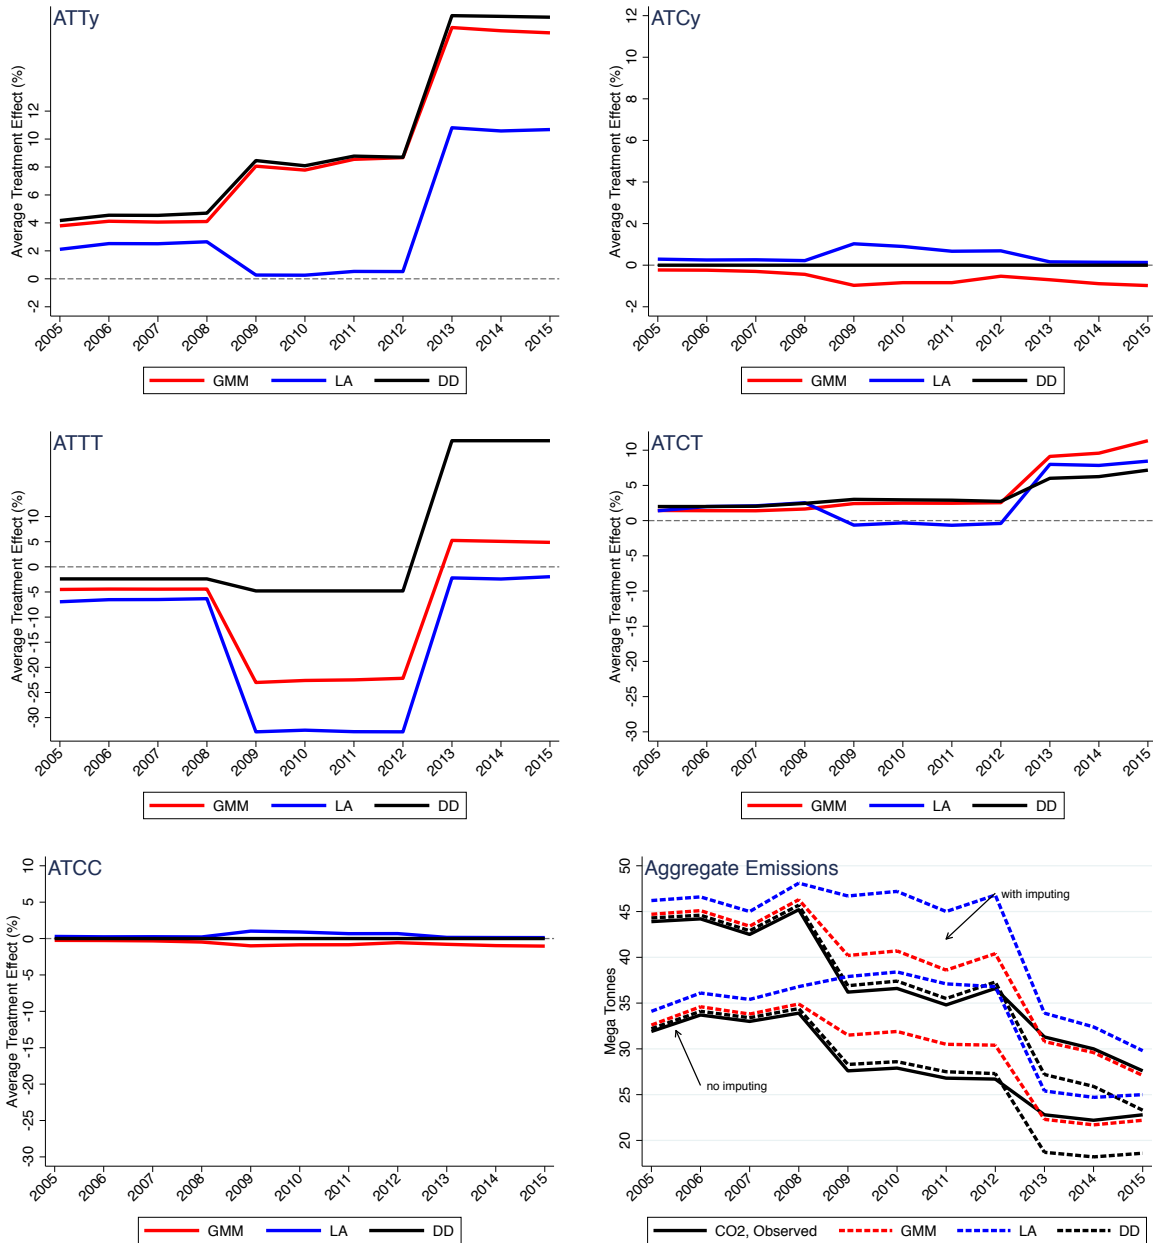
Notes: Figure presents the average effects on revenues for treated T (top left) and control C (top right) firms, excluding foreign imports, and on plant-level emissions for TT plants (middle left), CT plants (middle right) and CC plants (bottom left) generated for the domestic market (using plants in the EACEI surveys). Bottom right panel presents observed total CO₂ emissions generated for the domestic market (solid black) along with counterfactual emissions computed using GMM, LA and DD methods. Lower trajectories include only plants observed in the EACEI survey, while top trajectories impute emissions for small missing plants.

Figure A.5: Robustness Check, Alternative Matching Procedure



Notes: Figure presents the average effects on revenues for treated T (top left) and control C (top right) firms, using an alternative matching procedure, and on plant-level emissions for TT plants (middle left), CT plants (middle right) and CC plants (bottom left) generated for the domestic market (using plants in the EACEI surveys). Bottom right panel presents observed total CO₂ emissions generated for the domestic market (solid black) along with counterfactual emissions computed using GMM, LA and DD methods. Lower trajectories include only plants observed in the EACEI survey, while top trajectories impute emissions for small missing plants.

Figure A.6: Robustness Check, Alternative Industry Definition



Notes: Figure presents the average effects on revenues for treated T (top left) and control C (top right) firms, defining industries at the 3 digit level, and on plant-level emissions for TT plants (middle left), CT plants (middle right) and CC plants (bottom left) generated for the domestic market (using plants in the EACEI surveys). Bottom right panel presents observed total CO₂ emissions generated for the domestic market (solid black) along with counterfactual emissions computed using GMM, LA and DD methods. Lower trajectories include only plants observed in the EACEI survey, while top trajectories impute emissions for small missing plants.

B Theoretical Appendix

B.1 Proof of the exact CES price index change

In this section, we derive the change in CES price indices between t and t_0 as described in (3) with the definition of Sato-Vartia weights. The proof builds upon Feenstra (1994).

Using the expressions for price indices (3) and consumers' optimal expenditures (2), we can show that, for any firm f in the set Ω_{ist} of firms from industry i and sector s :

$$\frac{y_{fist}}{\sum_{\ell \in \Omega_{ist}} y_{list}} = p_{fist}^{\frac{\rho}{\rho-1}} P_{ist}^{-\frac{\rho}{\rho-1}},$$

where the LHS represents the market share of firm f within its industry at time t . For any pair of two periods t_0 and t , we define the set of continuing firms as $\Omega_{ist}^* \equiv \Omega_{ist} \cap \Omega_{ist_0}$. We can write:

$$\frac{y_{fist}}{\underbrace{\sum_{\ell \in \Omega_{ist}^*} y_{list}}_{\equiv \vartheta_{fist}}} \frac{\sum_{\ell \in \Omega_{ist}^*} y_{list}}{\underbrace{\sum_{\ell \in \Omega_{ist}} y_{list}}_{\equiv \lambda_{ist}}} = p_{fist}^{\frac{\rho}{\rho-1}} P_{ist}^{-\frac{\rho}{\rho-1}}.$$

In log differences between periods t_0 and t , we obtain

$$\log \frac{\vartheta_{fist}}{\vartheta_{fist_0}} + \log \frac{\lambda_{ist}}{\lambda_{ist_0}} = \frac{\rho}{\rho-1} \left(\log \frac{p_{fist}}{p_{fist_0}} - \log \frac{P_{ist}}{P_{ist_0}} \right),$$

which we rearrange and multiply on both sides by $\vartheta_{fist} - \vartheta_{fist_0}$ to get:

$$\frac{\rho-1}{\rho} (\vartheta_{fist} - \vartheta_{fist_0}) = \frac{\vartheta_{fist} - \vartheta_{fist_0}}{\log \frac{\vartheta_{fist}}{\vartheta_{fist_0}}} \log \frac{p_{fist}}{p_{fist_0}} - \frac{\vartheta_{fist} - \vartheta_{fist_0}}{\log \frac{\vartheta_{fist}}{\vartheta_{fist_0}}} \log \frac{P_{ist}}{P_{ist_0}} + \frac{1-\rho}{\rho} \frac{\vartheta_{fist} - \vartheta_{fist_0}}{\log \frac{\vartheta_{fist}}{\vartheta_{fist_0}}} \log \frac{\lambda_{ist}}{\lambda_{ist_0}}.$$

Summing over firms from the continuing set Ω_{ist}^* , we obtain:

$$\left(\sum_{\ell \in \Omega_{ist}^*} \frac{\vartheta_{list} - \vartheta_{list_0}}{\log \frac{\vartheta_{list}}{\vartheta_{list_0}}} \right) \log \frac{P_{ist}}{P_{ist_0}} = \sum_{\ell \in \Omega_{ist}^*} \frac{\vartheta_{fist} - \vartheta_{fist_0}}{\log \frac{\vartheta_{fist}}{\vartheta_{fist_0}}} \log \frac{p_{fist}}{p_{fist_0}} + \frac{1-\rho}{\rho} \left(\sum_{\ell \in \Omega_{ist}^*} \frac{\vartheta_{list} - \vartheta_{list_0}}{\log \frac{\vartheta_{list}}{\vartheta_{list_0}}} \right) \log \frac{\lambda_{ist}}{\lambda_{ist_0}},$$

which yields the first part of (11) and the definition of firm-specific Sato-Vartia weights:

$$\phi_{fist} \equiv \frac{\vartheta_{fist} - \vartheta_{fist_0}}{\log \vartheta_{fist} - \log \vartheta_{fist_0}} \bigg/ \sum_{\ell \in \Omega_{ist}^*} \frac{\vartheta_{list} - \vartheta_{list_0}}{\log \vartheta_{list} - \log \vartheta_{list_0}}.$$

Similarly, using (3) and (2) yields, for any industry i in the set Υ_s of industries from sector s :

$$\Theta_{ist} = P_{ist}^{\frac{\nu}{\nu-1}} \Psi_{st}^{-\frac{\nu}{\nu-1}}.$$

Taking the log differences between t_0 and t yields:

$$\frac{\log \frac{P_{ist}}{P_{ist_0}} - \log \frac{\Psi_{st}}{\Psi_{st_0}}}{\log \frac{\Theta_{ist}}{\Theta_{ist_0}}} = \frac{1 - \nu}{\nu}.$$

Multiplying each side of the equation by $\Theta_{ist} - \Theta_{ist_0}$ gives:

$$\frac{\Theta_{ist} - \Theta_{ist_0}}{\log \frac{\Theta_{ist}}{\Theta_{ist_0}}} \log \frac{P_{ist}}{P_{ist_0}} - \frac{\Theta_{ist} - \Theta_{ist_0}}{\log \frac{\Theta_{ist}}{\Theta_{ist_0}}} \log \frac{\Psi_{st}}{\Psi_{st_0}} = \frac{1 - \nu}{\nu} (\Theta_{ist} - \Theta_{ist_0}).$$

Summing over industries that belong to the stable set Υ_s from sector s , we obtain:

$$\sum_{i \in \Upsilon_s} \frac{\Theta_{ist} - \Theta_{ist_0}}{\log \frac{\Theta_{ist}}{\Theta_{ist_0}}} \log \frac{P_{ist}}{P_{ist_0}} = \left(\sum_{h \in \Upsilon_s} \frac{\Theta_{hst} - \Theta_{hst_0}}{\log \frac{\Theta_{hst}}{\Theta_{hst_0}}} \right) \log \frac{\Psi_{st}}{\Psi_{st_0}}$$

which yields to the second part of (11) and the definition of industry-specific Sato-Vartia weights:

$$\Phi_{ist} \equiv \frac{\Theta_{ist} - \Theta_{ist_0}}{\log \Theta_{ist} - \log \Theta_{ist_0}} \bigg/ \sum_{h \in \Upsilon_s} \frac{\Theta_{hst} - \Theta_{hst_0}}{\log \Theta_{hst} - \log \Theta_{hst_0}}.$$

B.2 Analytical Results on the True Treatment Effects for Single-Plant Firms

The average treatment effect can be expressed as the weighted average of sector-specific ATX_{st}^v 's as follows:

$$ATX_t^v = \sum_s \frac{N_{st}^X}{N_t^X} ATX_{st}^v, \text{ for } X \in \{T, C\} \text{ and } v \in \{y, Z\} \quad (\text{B.1})$$

where N_{st}^X and N_t^X denote the number of firms in group X at time t in sector s and the whole economy, respectively. Given the Cobb-Douglas assumption with respect to across-sector aggregation in the utility function, regulation in any given sector s has no effect on the $ATX_{s't}^v$ for any other sector s' . In order to sign ATX_t^v , it is thus sufficient to sign ATX_{st}^v , for any given s .

In this section, we demonstrate that:

Proposition 1. *When all firms operate a single plant,*

- i/ The average treatment effect on regulated (unregulated) firms' revenues is positive (negative) if and only if $\tau_t < 0$;*
- ii/ The average treatment effect on regulated firms' emissions is negative if $\tau_t > 0$, but can be either positive or negative if $\tau_t < 0$, depending on the aggregate market share*

of regulated firms. The average treatment effect on unregulated firms' emissions is positive if and only if $\tau_t > 0$;

iii/ The treatment effect on aggregate emissions is negative if and only if $\mu_t^z > 0$.

Average Treatment Effects for Revenues. Developing ATX_{st}^y first, we have

$$\begin{aligned} ATX_{st}^y &= \frac{1}{N_{st}^X} \sum_{i \in \Upsilon_s} \sum_{f \in \Omega_{ist}^X} \ln \frac{y_{fist}}{y'_{fist}} \\ &= \frac{1}{N_{st}^X} \sum_{i \in \Upsilon_s} \sum_{f \in \Omega_{ist}^X} \left(\ln \frac{\theta_{fist}}{\theta'_{fist}} + \ln \frac{\Theta_{ist}}{\Theta'_{ist}} \right), \end{aligned} \quad (\text{B.2})$$

where θ_{fist} and Θ_{ist} denote the within-industry market share of firm f and the within-sector market share of industry i at time t , respectively. The terms θ'_{fist} and Θ'_{ist} represent counterfactual market shares in the unregulated equilibrium.

Using (3) and (2), we rewrite

$$\frac{\theta_{fist}}{\theta'_{fist}} = \left(\frac{p_{fist}}{P_{ist}} \right)^{\frac{\rho}{\rho-1}} / \left(\frac{p'_{fist}}{P'_{ist}} \right)^{\frac{\rho}{\rho-1}} \quad (\text{B.3})$$

We relate observed (regulated) price to counterfactual price as follows: $p_{fist} = p'_{fist} e^{\tau_t R_{fist}}$. Substituting in the pricing rules yields

$$\frac{\theta_{fist}}{\theta'_{fist}} = \frac{e^{\frac{\rho \tau_t}{\rho-1} R_{fist}} \left[\sum_{j \in \Omega_t^T} (p'_{jist})^{\frac{\rho}{\rho-1}} + \sum_{k \in \Omega_t^C} (p'_{kist})^{\frac{\rho}{\rho-1}} \right]}{\sum_{j \in \Omega_t^T} (p'_{jist})^{\frac{\rho}{\rho-1}} e^{\frac{\rho \tau_t}{\rho-1} R_{jist}} + \sum_{k \in \Omega_t^C} (p'_{kist})^{\frac{\rho}{\rho-1}}} \quad (\text{B.4})$$

Taking logs and simplifying yields

$$\ln \frac{\theta_{fist}}{\theta'_{fist}} = \ln e^{\frac{\rho \tau_t}{\rho-1} R_{fist}} - \ln \left[\sum_{k \in \Omega_{ist}} \theta'_{kist} e^{\frac{\rho \tau_t}{\rho-1} R_{kist}} \right] \quad (\text{B.5})$$

Similarly, we develop

$$\ln \frac{\Theta_{ist}}{\Theta'_{ist}} = \frac{\nu(1-\rho)}{(1-\nu)\rho} \ln \left[\sum_{k \in \Omega_{ist}} \theta'_{kist} e^{\frac{\rho \tau_t}{\rho-1} R_{kist}} \right] - \ln \sum_{m \in \Upsilon_s} \Theta'_{mst} \left[\sum_{k \in \Omega_{mst}} \theta'_{kmst} e^{\frac{\rho \tau_t}{\rho-1} R_{kist}} \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}} \quad (\text{B.6})$$

Adding the two terms yields

$$\begin{aligned} \ln \frac{y_{fist}}{y'_{fist}} &= \ln e^{\frac{\rho}{\rho-1}\tau_t R_{fist}} - \frac{\rho-\nu}{(1-\nu)\rho} \ln \left[\sum_{k \in \Omega_{ist}} \theta'_{kist} e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} \right] \\ &\quad - \ln \sum_{m \in \Upsilon_s} \Theta'_{mst} \left[\sum_{k \in \Omega_{mst}} \theta'_{kmst} e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}}. \end{aligned} \quad (\text{B.7})$$

Plugging this expression into B.2 yields for treated firms

$$\begin{aligned} ATT_{st}^y &= \frac{\rho\tau_t}{\rho-1} \overline{R}_{st} - \frac{\rho-\nu}{(1-\nu)\rho} \sum_{i \in \Upsilon_s} \frac{N_{ist}^T}{N_{st}^T} \ln \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1 \right] \\ &\quad - \ln \sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1 \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}} \end{aligned} \quad (\text{B.8})$$

where $\overline{R}_{st} \equiv \frac{1}{N_{st}^T} \sum_{i \in \Upsilon_{st}} \sum_{f \in \Omega_{ist}^T} R_{fist}$, the average regulation value among regulated firms in the sector.

For single-plant firms, we set $R_{fist} = 1$ for all regulated plants. Equation B.8 simplifies to

$$\begin{aligned} ATT_{st}^y &= \ln e^{\frac{\rho\tau_t}{\rho-1}} - \frac{\rho-\nu}{(1-\nu)\rho} \sum_{i \in \Upsilon_s} \frac{N_{ist}^T}{N_{st}^T} \ln \left[\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1 \right] - \ln \sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1 \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}} \\ &= \underbrace{\frac{\rho-\nu}{\rho(1-\nu)} \sum_{i \in \Upsilon_s} \frac{N_{ist}^T}{N_{st}^T} \ln \left[\frac{e^{\frac{\rho\tau_t}{\rho-1}}}{\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1} \right]}_{>0} \underbrace{- \ln \left[\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1 \right]}_{-sign(\tau_t)} \underbrace{- \ln \left[\sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\frac{\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1}{e^{\frac{\rho\tau_t}{\rho-1}}} \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}} \right]}_{-sign(\tau_t)} \end{aligned} \quad (\text{B.9})$$

where $\zeta'_{ist} \equiv \sum_{k \in \Omega_{ist}^T} \theta'_{kist}$, the combined market share within industry i of regulated firms. Hence, we find that $ATT_{st}^y > 0 \iff \tau_t < 0$. If $\mu_t^e = \mu_t^z$, then $\tau_t = 0$ and $ATT_{st}^y = 0$. Iff $\mu_t^e > \mu_t^z$ then $\tau_t < 0$ and $ATT_{st}^y > 0$. We find that ATT_{st}^y is monotonically increasing in μ_t^e .

For control firms, using B.7, we can write

$$\begin{aligned} ATC_{st}^y &= - \left\{ \underbrace{\frac{\rho-\nu}{(1-\nu)\rho} \sum_{i \in \Upsilon_s} \frac{N_{ist}^C}{N_{st}^C} \ln \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1 \right]}_{>0} \right. \\ &\quad \left. + \underbrace{\ln \sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1 \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}}}_{-sign(\tau_t)} \right\}. \end{aligned} \quad (\text{B.10})$$

Hence, we find that $ATC_{st}^y > 0 \iff \tau_t > 0$. Aggregating over sectors yields $sign(AT T_t^y) = -sign(ATC_t^y) = -sign(\tau_t)$.

Average Treatment Effect for Emissions. Second, developing ATX_{st}^z yields

$$ATX_{st}^z = \frac{1}{N_{st}^X} \sum_{i \in \Upsilon_s} \sum_{f \in \Omega_{ist}^X} \ln \frac{Z_{fist}}{Z'_{fist}} = \frac{1}{N_{st}^X} \sum_{i \in \Upsilon_s} \sum_{f \in \Omega_{ist}^X} \left[\ln e^{-\mu_t^Z R_{fist}} + \ln \left(\frac{y_{fist}}{y'_{fist}} \right) \right] \quad (\text{B.11})$$

Using B.7 yields

$$\begin{aligned} ATX_{st}^z &= \frac{1}{N_{st}^X} \sum_{i \in \Upsilon_{st}} \sum_{f \in \Omega_{ist}^X} \ln e^{(-\mu^Z + \frac{\rho\tau_t}{\rho-1})R_{fist}} \\ &\quad - \frac{\rho - \nu}{(1 - \nu)\rho} \sum_{i \in \Upsilon_s} \frac{N_{ist}^X}{N_{st}^X} \ln \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1 \right] \\ &\quad - \ln \sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1 \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}}. \end{aligned} \quad (\text{B.12})$$

For single-plant firms, we set $R_{fist} = 1$ for all regulated plants. For treated firms, equation B.12 simplifies to

$$ATT_{st}^z = \frac{\rho - \nu}{\rho(1 - \nu)} \sum_{i \in \Upsilon_s} \frac{N_{ist}^T}{N_{st}^T} \ln \left[\frac{e^{(-\mu^Z + \frac{\rho\tau_t}{\rho-1})}}{\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1} \right] - \ln \left[\sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\frac{\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1}{e^{(-\mu^Z + \frac{\rho\tau_t}{\rho-1})}} \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}} \right] \quad (\text{B.13})$$

We obtain

$$ATT_{st}^z > 0 \iff \frac{e^{(-\mu^Z + \frac{\rho\tau_t}{\rho-1})}}{\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1} > 1.$$

As a result, if $\tau_t > 0$, then $\frac{e^{(-\mu^Z + \frac{\rho\tau_t}{\rho-1})}}{\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1} < 1$ and we conclude that $ATT_{st}^z < 0$. By contrast,

if $\tau_t < 0$, then the sign of the term $\frac{e^{(-\mu^Z + \frac{\rho\tau_t}{\rho-1})}}{\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1} - 1$ depends on ζ'_{ist} , the aggregate market

share of regulated firms in the unregulated counterfactual. The sign of ATT_{st}^z is thus ambiguous if $\tau_t < 0$. Furthermore, if $\mu_t^e = \mu_t^z$ then $\tau_t = 0$ and $ATT_{st}^z = -\mu_t^z$. Given that $ATT_{st}^z = -\mu_t^z + ATT_{st}^y$, then ATT_{st}^z is monotonically increasing in μ_t^e .

For control firms, equation B.12 simplifies to

$$\begin{aligned}
ATC_{st}^z &= - \left\{ \frac{\rho - \nu}{(1 - \nu)\rho} \sum_{i \in \Upsilon_s} \frac{N_{ist}^C}{N_{st}^C} \ln \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1 \right] \right. \\
&\quad \left. + \ln \sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1 \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}} \right\} \quad (B.14)
\end{aligned}$$

which is positive if and only if $\tau_t > 0$.

Aggregate effects on emissions. We define ATZ_t as the log difference between the sum of observed and counterfactual firm-level emissions. Using (13), we express

$$ATZ_t \equiv \ln \left[\frac{\sum_{f \in \Omega_t} Z_{fist}}{\sum_{f \in \Omega_t} Z'_{fist}} \right] = \ln \left[\frac{\sum_{f \in \Omega_t} \frac{Z_{fist}}{y_{fist}} y_{fist}}{\sum_{f \in \Omega_t} \frac{Z'_{fist}}{y'_{fist}} y'_{fist}} \right] = \ln \left[\frac{\sum_{f \in \Omega_t} e^{-\mu_t^Z} R_{fist} y'_{fist} \frac{y_{fist}}{y'_{fist}}}{\sum_{f \in \Omega_t} y'_{fist}} \right] \quad (B.15)$$

Using (B.7), we obtain

$$ATZ_t = \ln \left[\sum_s a_{st} \left(\frac{\sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\sum_{k \in \Omega_{ist}} \theta'_{kist} e^{(\frac{\rho\tau_t}{\rho-1} - \mu_t^Z) R_{kist}} \right] \left[\sum_{k \in \Omega_{ist}} \theta'_{kist} e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} \right]^{-\frac{(\rho-\nu)}{(1-\nu)\rho}}}{\sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\sum_{k \in \Omega_{ist}} \theta'_{kist} e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}}} \right) \right] \quad (B.16)$$

Assuming single-plant firms, this expression simplifies to

$$ATZ_t = \ln \left[\sum_s a_{st} \left(\frac{\sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\zeta'_{ist} \left(e^{(\frac{\rho\tau_t}{\rho-1} - \mu_t^Z) - 1} \right) + 1 \right] \left[\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1} - 1} \right) + 1 \right]^{-\frac{(\rho-\nu)}{(1-\nu)\rho}}}{\sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1} - 1} \right) + 1 \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}}} \right) \right] \quad (B.17)$$

From this expression, we see that $ATZ_t = 0$ if $\mu_t^z = 0$. Furthermore, we have $ATZ_t < 0$ for any $\mu_t^z > 0$.

B.3 Analytical Results on the Bias from the DD Estimator for Single-Plant Firms

In this appendix, we demonstrate that

Proposition 2. *When all firms operate a single plant, the DD estimator*

- i/ exaggerates the magnitude of ATT_t^y ,
- ii/ exaggerates the magnitude of ATT_t^Z if $\tau_t > 0$, but is biased upwards if $\tau_t < 0$. In the latter case, the $E \left[\widehat{ATT_t^{Z,DD}} \right]$ could be biased towards zero or away from zero, and could even yield the opposite sign from the true effect.
- iii/ exaggerates the magnitude of ATZ_t if $\tau_t + \epsilon > 0$, but is biased upwards if $\tau_t + \epsilon < 0$. In the latter case, the $E \left[\widehat{ATZ_t^{DD}} \right]$ could yield the opposite sign from the true effect.

Bias in ATT^y from the DD estimator. After estimating the DD regression (27), we define the ATT in revenues as

$$\widehat{ATT_t^{y,DD}} = \frac{1}{N_t^T} \sum_{f \in \Omega_t^T} \widehat{\beta_{0,t}^{y,DD}} R_{fist} \quad (\text{B.18})$$

where $\widehat{\beta_{0,t}^{y,DD}}$ is the estimated coefficient. In the case that all firms own a single plant, $\widehat{ATT_t^{y,DD}} = \widehat{\beta_{0,t}^{y,DD}}$. We thus compute the bias in the DD estimate of the ATT on revenues for the case of single-plant firms as

$$Bias_{st}^{y,DD} = E \left[\widehat{ATT_t^{y,DD}} \right] - ATT_t^y = \frac{\rho \tau_t}{\rho - 1} - ATT_t^y \quad (\text{B.19})$$

where $E \left[\widehat{\beta_{0,t}^{y,DD}} \right] = \frac{\rho \tau_t}{\rho - 1}$ if regulation is orthogonal to unobserved productivity growth.

Using (B.9)

$$Bias_{st}^{y,DD} = \frac{\rho - \nu}{(1 - \nu)\rho} \sum_{i \in \Upsilon_s} \frac{N_{ist}^T}{N_{st}^T} \ln \left[\zeta'_{ist} \left(e^{\frac{\rho \tau_t}{\rho - 1}} - 1 \right) + 1 \right] + \ln \sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\zeta'_{ist} \left(e^{\frac{\rho \tau_t}{\rho - 1}} - 1 \right) + 1 \right]^{\frac{\nu(\rho - 1)}{(\nu - 1)\rho}} \quad (\text{B.20})$$

We obtain $Bias_{st}^{y,DD} > 0 \iff \tau_t < 0$. By proposition 1, this implies that $Bias_{st}^{y,DD} > 0 \iff ATT_{st}^y > 0$, i.e., the DD estimator is biased upward in magnitude for a given sector. Aggregating over sectors implies that the DD estimate of ATT_t^y is biased upward in magnitude.

Bias in ATT^Z from the DD estimator. We define the ATT in emissions computed via the DD method as

$$\widehat{ATT_t^{Z,DD}} = \frac{1}{N_t^T} \sum_{f \in \Omega_t^T} \widehat{\beta_{0,t}^{Z,DD}} R_{fist} \quad (\text{B.21})$$

where $\widehat{\beta_{0,t}^{Z,DD}}$ is the coefficient estimated from the DD regression (27). In the case that all firms own a single plant, $\widehat{ATT_t^{Z,DD}} = \widehat{\beta_{0,t}^{Z,DD}}$. We thus compute the bias in the DD estimate

of the ATT on emissions for the case of single-plant firms as

$$Bias_t^{Z,DD} = E \left[\widehat{ATT}_t^{Z,DD} \right] - ATT_t^Z = -\mu_t^Z + \frac{\rho\tau_t}{\rho-1} - ATT_t^Z \quad (\text{B.22})$$

where $E \left[\widehat{\beta}_{0,t}^{Z,DD} \right] = -\mu_t^Z + \frac{\rho\tau_t}{\rho-1}$ if regulation is orthogonal to unobserved productivity growth.

Using (B.13), we express the bias for a given sector s as

$$Bias_{st}^{Z,DD} = \frac{\rho-\nu}{(1-\nu)\rho} \sum_{i \in \Upsilon_s} \frac{N_{ist}^T}{N_{st}^T} \ln \left[\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1 \right] + \ln \sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1 \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}} \quad (\text{B.23})$$

Comparing (B.20) and (B.23) reveals that $Bias_{st}^{Z,DD} = Bias_{st}^{y,DD}$. Hence, $Bias_{st}^{Z,DD} > 0 \iff \tau_t < 0$. By proposition 1, this implies that if $\tau_t > 0$, then $E \left[\widehat{ATT}_t^{Z,DD} \right] < ATT_t^Z < 0$; hence, the DD estimator overstates the negative treatment effect on emissions.

By contrast, if $\tau_t < 0$, then $Bias_{st}^{Z,DD} > 0$, but since the sign of ATT_{st}^Z is ambiguous in this case, given proposition 1, $\widehat{ATT}_t^{Z,DD}$ could be biased towards 0 (if $ATT_{st}^Z < 0$) or away from 0 (if $ATT_{st}^Z > 0$). When $ATT_{st}^Z < 0$, $\widehat{ATT}_t^{Z,DD}$ could be biased upwards so much that it would take the opposite sign of the true ATT on emissions in expectation.

Bias in ATZ from the DD estimator. The DD estimate of the aggregate effect on emissions can be written as

$$\widehat{ATZ}_t^{DD} = \ln \left[\frac{\sum_{f \in \Omega_t} Z_{fist}}{\sum_{f \in \Omega_t} Z_{fist} e^{-\widehat{\beta}_{0,t}^{Z,DD} R_{fist}}} \right] \quad (\text{B.24})$$

where $\widehat{\beta}_{0,t}^{Z,DD}$ is the coefficient estimated from the DD regression (27). For single-plant firms, we can rewrite it as

$$\widehat{ATZ}_t^{DD} = \ln \left[\sum_s a_s \left(\frac{\sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\zeta'_{ist} \left(e^{\left(\frac{\rho\tau_t}{\rho-1} - \mu_t^Z \right)} - 1 \right) + 1 \right] \left[\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1 \right]^{-\frac{(\rho-\nu)}{(1-\nu)\rho}}}{\sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1 \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}}} \right) \right] - \ln \left[\sum_s a_s \left(\frac{\sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\zeta'_{ist} \left(e^{\left(\frac{\rho\tau_t}{\rho-1} - \mu_t^Z - \widehat{\beta}_{0,t}^{Z,DD} \right)} - 1 \right) + 1 \right] \left[\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1 \right]^{-\frac{(\rho-\nu)}{(1-\nu)\rho}}}{\sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\zeta'_{ist} \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + 1 \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}}} \right) \right] \quad (\text{B.25})$$

Due to Jensen inequality, we have

$$E[\widehat{ATZ^{DD}}(\widehat{\beta_{0,t}^{Z,DD}})] - ATZ_t \leq \widehat{ATZ^{DD}}(E[\widehat{\beta_{0,t}^{Z,DD}}]) - ATZ_t, \quad (\text{B.26})$$

where the left hand side of the inequality corresponds to the bias from the DD estimator and the right hand side to the approximation of the bias that we can compute.

Since $E\left[\widehat{\beta_{0,t}^{Z,DD}}\right] = -\mu_t^Z + \frac{\rho\tau_t}{\rho-1}$ if regulation is orthogonal to unobserved productivity growth, we have $\widehat{ATZ^{DD}}(E[\widehat{\beta_{0,t}^{Z,DD}}]) - ATZ_t = 0$ when $\tau_t = 0$. Furthermore, we obtain $\text{sign}(\widehat{ATZ^{DD}}(E[\widehat{\beta_{0,t}^{Z,DD}}]) - ATZ_t) = -\text{sign}(\tau_t)$.

As a result, when $\tau_t + \epsilon > 0$ (due to Jensen inequality), the bias from the DD estimator is negative, thereby implying that the DD estimator overstates the fall in aggregate emissions. When $\tau_t + \epsilon < 0$, the DD estimator underestimates the fall in aggregate emissions, and might even predict the wrong sign of the effect.

B.4 Analytical Results on the Bias from the LA Estimator for Single-Plant Firms

LA model as a local approximation. We first show that estimation equation (29) follows from taking a local approximation of changes around an equilibrium set at time t_0 .

For any vector of infinitesimal changes in firm-level prices relative to the equilibrium obtained at time t_0 , noted $\{dp_{fist_0}\}_{f \in \Omega_{ist_0}}$, the change in the industry price index P_{ist} is given by:

$$dP_{ist_0} = \sum_{f \in \Omega_{ist_0}} \frac{\partial P_{ist}}{\partial p_{fist} |_{t=t_0}} dp_{fist_0}. \quad (\text{B.27})$$

Given (3), we have that, at any period t :

$$\frac{\partial P_{ist}}{\partial p_{fist}} = \left(\sum_{f \in \Omega_{ist}} (p_{fist})^{\frac{\rho}{\rho-1}} \right)^{\frac{\rho-1}{\rho}-1} (p_{fist})^{\frac{\rho}{\rho-1}-1} = \theta_{fist} \frac{P_{ist}}{p_{fist}},$$

which we can plug in (B.27) to obtain:

$$dP_{ist_0} = \sum_{f \in \Omega_{ist_0}} \theta_{fist_0} dp_{fist_0}. \quad (\text{B.28})$$

As a result, if we approximate small variation in firms' prices with infinitesimal changes around t_0 equilibrium, we obtain $\Delta P_{ist_0} \approx \sum_{f \in \Omega_{ist_0}} \theta_{fist_0} \Delta p_{fist_0}$ for the change in industry i 's price index. We could similarly build $\Delta \Psi_{st_0} \approx \sum_{i \in \Upsilon_{st_0}} \Theta_{fist_0} \Delta P_{ist_0}$ for the change in sector s 's price index, but controlling for sector-time fixed effects absorbs these spillovers.

In the rest of the section, we demonstrate that

Proposition 3. *When all firms operate a single plant, the LA estimator*

i/ yields unbiased estimates of the average treatments on the regulated and unregulated firms' revenues and emissions, as well as of the effect on aggregate emissions, if there is no entry and exit of firms and if market shares remain at their pre-regulation values;

ii/ otherwise, it yields estimates of average and aggregate effects that can be biased up or down and can even have the opposite sign from the true effects.

LA model under the condition that the approximation is valid. Under the joint condition that

$$\vartheta_{fist_0} \approx \vartheta_{fist} \approx \vartheta'_{fist}, \quad \Theta_{ist_0} \approx \Theta_{ist} \approx \Theta'_{ist}, \quad \Delta\lambda_{it} \approx \Delta\lambda'_{it} \approx 0 \quad (\text{B.29})$$

then denoting $\beta_{0,t}^y = \frac{\rho\tau_t}{\rho-1}$ and $\beta_{0,t}^z = -\mu_t^Z + \frac{\rho\tau_t}{\rho-1}$, (12) can be written as

$$\begin{aligned} \Delta v_{fist} &= \beta_{0,t}^v R_{fist} + \frac{(\nu - \rho)\tau_t}{(\nu - 1)(1 - \rho)} \sum_{k \in \Omega_{ist}} \theta_{kist_0} R_{kist} + \frac{\nu\tau}{1 - \nu} \sum_{m \in \Upsilon_s} \Theta_{mst_0} \sum_{\ell \in \Omega_{mst}} \theta_{\ell mst_0} R_{\ell mst} \\ &+ \frac{(\nu - \rho)}{(\nu - 1)(\rho - 1)} \sum_{k \in \Omega_{ist}} \theta_{kist_0} \Delta A_{kist} + \frac{\nu}{\nu - 1} \sum_{m \in \Upsilon_s} \Theta_{mst_0} \sum_{\ell \in \Omega_{mst}^*} \theta_{\ell mst_0} \Delta A_{\ell mst} + \Delta Y_{st} + \frac{\rho}{1 - \rho} \Delta A_{fist} \end{aligned}$$

where we have substituted base-year market shares θ s for firm-level Sato-Vartia weights ϕ s, and base-year industry-level market shares Θ s for industry-level Sato-Vartia weights Φ s.

In this case, the last line is orthogonal to firm-level regulation R_{fist} and industry-level regulation $\sum_{k \in \Omega_{ist}} \theta_{kist_0} R_{kist}$. Hence, the OLS regression (29) identifies $\beta_{0,t}^y$, $\beta_{0,t}^z$ and $\beta_{2,t}^y$, i.e., $E \left[\widehat{\beta_{0,t}^{y,LA}} \right] = \frac{\rho\tau_t}{\rho-1}$, $E \left[\widehat{\beta_{0,t}^{z,LA}} \right] = -\mu_t^Z + \frac{\rho\tau_t}{\rho-1}$ and $E \left[\widehat{\beta_{2,t}^{y,LA}} \right] = \frac{(\nu-\rho)\tau_t}{(\nu-1)(1-\rho)}$, and firm-level estimated treatment effects coincide with the true firm-level treatment effects in expectation:

$$\begin{aligned} E \left[\ln \left(\frac{v_{fist}}{v'_{fist}} \right) \right] &= \beta_{0,t}^v R_{fist} + \frac{(\nu - \rho)\tau_t}{(\nu - 1)(1 - \rho)} \sum_{k \in \Omega_{ist_0}} \theta_{kist_0} R_{kist} \\ &- \left(\frac{\rho\tau_t}{\rho - 1} + \frac{(\nu - \rho)\tau_t}{(\nu - 1)(1 - \rho)} \right) \sum_{m \in \Upsilon_s} \Theta_{mst_0} \sum_{k \in \Omega_{mst_0}} \theta_{kmst_0} R_{kmst} \\ &= \ln \left(\frac{v_{fist}}{v'_{fist}} \right). \end{aligned}$$

Since the LA model yields unbiased estimates of firm-level treatment effects, in expectation,

aggregating over firms, we have $E \left[\widehat{ATT}_t^{v,LA} \right] = ATT_t^v$ and $E \left[\widehat{ATZ}_t^{LA} \right] = ATZ_t$.

LA model under the condition that the approximation is not valid. Conversely, if condition (B.29) does not hold, then the bias in the estimate of average treatment effects on outcome $v \in \{y, Z\}$ for firms in group $X \in \{T, C\}$ in sector s for the LA model relative to (B.9) can be computed as:

$$\begin{aligned} Bias_{st}^{v,LA,X} &= E \left[\widehat{\beta}_{0,t}^{v,LA} \right] \overline{R}_{st} \mathbb{1}(X = T) + E \left[\widehat{\beta}_{2,t}^{y,LA} \right] \sum_{i \in \Upsilon_s} \frac{N_{ist}^X}{N_{st}^X} \zeta_{ist_0} - E \left[\left(\widehat{\beta}_{0,t}^{y,LA} + \widehat{\beta}_{2,t}^{y,LA} \right) \right] \zeta_{st_0} \\ &- \left\{ \beta_{0,t}^v \overline{R}_{st} \mathbb{1}(X = T) + \frac{\rho - \nu}{\rho(1 - \nu)} \sum_{i \in \Upsilon_s} \frac{N_{ist}^X}{N_{st}^X} \ln \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho \tau_t}{\rho - 1} R_{kist}} - 1 \right) + 1 \right] \right. \\ &\left. + \ln \left[\sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho \tau_t}{\rho - 1} R_{kist}} - 1 \right) + 1 \right]^{\frac{\nu(\rho - 1)}{(\nu - 1)\rho}} \right] \right\}, \end{aligned} \quad (B.30)$$

where $\widehat{\beta}_{0,t}^{v,LA}$ and $\widehat{\beta}_{2,t}^{y,LA}$ are computed from regression equation (29) and where $\zeta_{ist_0} \equiv \sum_{k \in \Omega_{ist}^T} \theta'_{kist_0} R_{kist}$ corresponds to the weighted market share of regulated firms in base-year t_0 in industry i , and $\zeta_{st_0} \equiv \sum_m \Theta_{mst} \zeta_{mst_0}$ its counterpart for sector s , and $\overline{R}_{st} \equiv \frac{1}{N_{st}^T} \sum_{k \in \Omega_{ist}^T} R_{kist}$ is the average regulation among regulated firms. In the single-plant case, $\overline{R}_{st} = 1$.

The bias can be decomposed as follows:

$$Bias_{st}^{v,LA,X} = \text{Reg-Bias}_{st}^{v,LA,X} + \text{Approx-Bias}_{st}^{LA,X} \quad (B.31)$$

with

$$\begin{aligned} \text{Reg-Bias}_{st}^{v,LA,X} &= \left(E \left[\widehat{\beta}_{0,t}^{v,LA} \right] - \beta_{0,t}^v \right) \overline{R}_{st} \mathbb{1}(X = T) + \left(E \left[\widehat{\beta}_{2,t}^{y,LA} \right] - \beta_{2,t}^y \right) \sum_{i \in \Upsilon_s} \frac{N_{ist}^X}{N_{st}^X} \zeta_{ist_0} \\ &- \left(E \left[\widehat{\beta}_{0,t}^{y,LA} \right] + E \left[\widehat{\beta}_{2,t}^{y,LA} \right] - (\beta_{0,t}^y + \beta_{2,t}^y) \right) \zeta_{st_0} \end{aligned} \quad (B.32)$$

and

$$\begin{aligned} \text{Approx-Bias}_{st}^{LA,X} &= \frac{\beta_{2,t}^y}{\beta_{0,t}^y} \sum_{i \in \Upsilon_s} \frac{N_{ist}^X}{N_{st}^X} \ln \left[\frac{e^{\beta_{0,t}^y \zeta_{ist_0}}}{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\beta_{0,t}^y R_{kist}} - 1 \right) + 1} \right] \\ &- \ln \left[\sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\frac{e^{\beta_{0,t}^y \zeta_{st_0}}}{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\beta_{0,t}^y R_{kist}} - 1 \right) + 1} \right]^{\left(\frac{\beta_{2,t}^y + \beta_{0,t}^y}{\beta_{0,t}^y} \right)} \right] \end{aligned} \quad (B.33)$$

The term $\text{Reg-Bias}_{st}^{v,LA,X}$ represents the bias in the estimated treatment effect resulting from using biased coefficients in the computation of firm-level treatment effects (30). This

bias arises because the error term in the regression (29) includes Sato-Vartia-weighted productivity growth at the industry level, which correlates with industry-average regulation density. The additional term $\text{Approx-Bias}_{st}^{LA,X}$ represents the bias resulting from using a linear approximation to the CES price index at t_0 .

In the single-plant case, $\text{Approx-Bias}_{st}^{LA,X}$ simplifies to

$$\text{Approx-Bias}_{st}^{LA,X} = \frac{\beta_{2,t}^y}{\beta_{0,t}^y} \sum_{i \in \mathcal{Y}_s} \frac{N_{ist}^X}{N_{st}^X} \ln \left[\frac{e^{\beta_{0,t}^y \zeta_{ist_0}}}{\zeta'_{ist} (e^{\beta_{0,t}^y} - 1) + 1} \right] - \ln \left[\sum_{i \in \mathcal{Y}_s} \Theta'_{ist} \left[\frac{e^{\beta_{0,t}^y \zeta_{st_0}}}{\zeta'_{ist} (e^{\beta_{0,t}^y} - 1) + 1} \right]^{\left(\frac{\beta_{2,t}^y + \beta_{0,t}^y}{\beta_{0,t}^y} \right)} \right] \quad (\text{B.34})$$

with $\zeta'_{ist} = \sum_{k \in \Omega_{ist}^T} \theta'_{kist}$, the market share of regulated firms in the counterfactual equilibrium at t . This term is positive if, for all industries i ,

$$e^{\beta_{0,t}^y \zeta_{ist_0}} - 1 < \zeta'_{ist} (e^{\beta_{0,t}^y} - 1) \quad (\text{B.35})$$

and

$$e^{\beta_{0,t}^y \zeta_{st_0}} - 1 < \zeta'_{ist} (e^{\beta_{0,t}^y} - 1). \quad (\text{B.36})$$

If $\zeta_{ist_0} = \zeta'_{ist} = \zeta_{st_0}$, then these conditions hold for any ζ_{st} . To see this, observe that the function $F(\beta_{0,t}^y, \zeta_{st}) \equiv e^{\beta_{0,t}^y \zeta_{st}} - 1 - \zeta_{st} (e^{\beta_{0,t}^y} - 1)$ has a global maximum at 0. Hence, conditions (B.35) and (B.36) hold, so $\text{Approx-Bias}_{st}^{LA}$ is positive, if $\zeta_{ist_0} = \zeta'_{ist} = \zeta_{st_0}$.

Generally, conditions (B.35) and (B.36) may or may not hold, depending on the values of ζ_{ist_0} , ζ'_{ist} , ζ_{st_0} . For example, suppose $\zeta_{ist_0} = \zeta'_{ist} = \zeta_{st_0} = .5$ and $\beta_{0,t}^y = 2$, then $e^{\beta_{0,t}^y \zeta_{ist_0}} - 1 - \zeta'_{ist} (e^{\beta_{0,t}^y} - 1) = -1.47 < 0$, but if $\zeta_{ist_0} = \zeta_{st_0} = .5$ and $\zeta'_{ist} = .1$ and $\beta_{0,t}^y = 2$, then $e^{\beta_{0,t}^y \zeta_{ist_0}} - 1 - \zeta'_{ist} (e^{\beta_{0,t}^y} - 1) = 1.08 > 0$. Hence, (B.35) may or may not hold, and $\text{Approx-Bias}_{st}^{LA}$ may be positive or negative, conditional on the same set of parameters. Thus, in the single plant case, we cannot sign $\text{Bias}_{st}^{v,LA,X}$ analytically.

The LA estimate of the aggregate effect on emissions can be written as

$$\widehat{ATZ}_t^{LA} = \ln \left[\frac{\sum_{f \in \Omega_t} Z_{fist}}{\sum_{f \in \Omega_t} Z_{fist} e^{-\left[\widehat{\beta}_{0,t}^{Z,LA} R_{fist} + \widehat{\beta}_{2,t}^{Z,LA} \zeta_{ist_0} - \left(\widehat{\beta}_{0,t}^{Z,LA} + \widehat{\beta}_{2,t}^{Z,LA} \right) \zeta_{st_0} \right]}} \right] \quad (\text{B.37})$$

where $\widehat{\beta}_{0,t}^{Z,LA}$ and $\widehat{\beta}_{2,t}^{Z,LA}$ are computed from regression model (29). As above, we can decompose the bias from the LA estimator into two components: the first component arises because the error term in this model includes Sato-Vartia-weighted productivity growth at the industry level, which generates an omitted variable bias, and the second component results from using a linear approximation to the CES price index at t_0 . As before, these components have an indeterminate sign given the structural parameters. Hence, we cannot sign the bias in the estimated ATZ from the LA model analytically.

B.5 Analytical Results on the True Firm-Level Treatment Effects for Multi-Plant Firms

In this section, we start by demonstrating that

Proposition 4. *When firms operate multiple plants,*

- i/ The average treatment effect on revenues for regulated firms could be positive or negative, whatever the sign of τ_t . The average treatment effect on revenues for unregulated firms is positive if and only if $\tau_t > 0$;*
- ii/ The average treatment effect on emissions for regulated firms could be positive or negative, whatever the sign of τ_t . The average treatment effect on emissions for unregulated firms is positive if and only if $\tau_t > 0$;*
- iii/ Aggregate emissions decreases if and only if $\mu_t^Z > 0$.*

Average Treatment Effect on Revenues. The general expression for ATT_{st}^y is given by (B.8) with a continuous measure of treatment, $R_{fist} \in [0, 1]$, which can be rewritten as

$$\begin{aligned}
 ATT_{st}^y &= \frac{\rho - \nu}{(1 - \nu)\rho} \sum_{i \in \Upsilon_s} \frac{N_{ist}^T}{N_{st}^T} \ln \left[\frac{e^{\frac{\rho\tau_t}{\rho-1} \overline{R}_{st}}}{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1} \right] \\
 &\quad - \ln \sum_{i \in \Upsilon_s} \theta'_{ist} \left[\frac{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1}{e^{\frac{\rho\tau_t}{\rho-1} \overline{R}_{st}}} \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}}
 \end{aligned} \tag{B.38}$$

whose sign depends on the following inequality:

$$e^{\frac{\rho\tau_t}{\rho-1} \overline{R}_{st}} - 1 \leq \sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right), \quad \forall i. \tag{B.39}$$

If the inequality in (B.39) can be signed, conditional on parameters, then so can the ATT_{st}^y and ATT_t^y . Conversely, if the direction of the inequality in (B.39) can vary for given parameters, then the ATT_{st}^y cannot be signed analytically. In this case, the sign of the ATT_{st}^y depends on the entire vector of regulation $\{R_{fist}\}$.

A simple example provides an illustration of the ambiguity of the sign in ATT_{st}^y . Suppose there is a single industry with two treated firms and several control firms. The regulation of the first regulated firm is $R_1 = 1$, and the regulation of the second regulated firm is $R_2 = \epsilon$. The ATT is written as:

$$ATT_t^y = \frac{\rho\tau_t}{\rho-1} \frac{1 + \epsilon}{2} - \ln \left[\theta'_1 \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + \theta'_2 \left(e^{\frac{\rho\tau_t}{\rho-1} \epsilon} - 1 \right) + 1 \right]$$

Suppose $\theta'_1 = \theta'_2 = 0.4$, $\frac{\rho\tau_t}{\rho-1} = 2$, and $\epsilon = 0.1$. Then the $ATT_t^y = -0.193$. But if $\epsilon = 0.4$, then $ATT_t^y = 0.002$. Now suppose $\theta'_1 = .7$ and $\theta'_2 = 0.25$, $\frac{\rho\tau_t}{\rho-1} = -1$, and $\epsilon = 0.1$. Then the

$ATT_t^y = 0.077$. But if $\epsilon = 0.7$, then $ATT_t^y = -0.010$. This example shows that whatever the sign of τ_t , the ATT can be positive or negative, depending on the entire vector of regulation.

The expression (B.10) for ATC_{st}^y is general, as it allows for continuous treatment $R_{fist} \in [0, 1]$. Hence, the result for control firms in the multi-plant case is immediate: $ATC_t^y > 0 \iff \tau_t > 0$.

Average Treatment Effect on Emissions. The general expression for ATT_{st}^Z is given by (B.12) for a continuous treatment $R_{fist} \in [0, 1]$, which can be rewritten as

$$ATT_{st}^Z = \frac{\rho - \nu}{(1 - \nu)\rho} \sum_{i \in \mathcal{Y}_s} \frac{N_{ist}^T}{N_{st}^T} \ln \left[\frac{e^{(-\mu_t^Z + \frac{\rho\tau_t}{\rho-1})\overline{R}_{st}}}{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1} \right] - \ln \sum_{i \in \mathcal{Y}_s} \theta'_{ist} \left[\frac{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1}{e^{(-\mu_t^Z + \frac{\rho\tau_t}{\rho-1})\overline{R}_{st}}} \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}} \quad (\text{B.40})$$

whose sign depends on the following inequality

$$e^{(-\mu_t^Z + \frac{\rho\tau_t}{\rho-1})\overline{R}_{st}} - 1 \leq \sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right), \quad \forall i. \quad (\text{B.41})$$

If the inequality in (B.41) can be signed, conditional on parameters, then so can the ATT_{st}^Z , and ATT_t^Z . Conversely, if the direction of the inequality in (B.41) can vary for given parameters, then the ATT_{st}^Z cannot be signed analytically, as its sign depends on the full vector of regulation.

A simple example provides an illustration of the ambiguity of the sign in ATT_{st}^Z . Suppose there is a single industry with two treated firms and several control firms. The regulation of the first regulated firm is $R_1 = 1$, and the regulation of the second regulated firm is $R_2 = \epsilon$. The ATT^Z is written as:

$$ATT_t^Z = \left(-\mu_t^Z + \frac{\rho\tau_t}{\rho-1} \right) \left(\frac{1 + \epsilon}{2} \right) - \ln \left[\theta'_1 \left(e^{\frac{\rho\tau_t}{\rho-1}} - 1 \right) + \theta'_2 \left(e^{\frac{\rho\tau_t}{\rho-1}\epsilon} - 1 \right) + 1 \right]$$

Suppose $\theta'_1 = \theta'_2 = 0.4$, $\frac{\rho\tau_t}{\rho-1} = 2$, and $\mu_t^Z = .05$. If $\epsilon = 0.1$, then $ATT_t^Z = -0.221$, whereas, if $\epsilon = 0.5$, $ATT_t^Z = 0.017$. Now suppose $\theta'_1 = .7$ and $\theta'_2 = 0.25$, $\frac{\rho\tau_t}{\rho-1} = -1$, and $\mu_t^Z = .05$. If $\epsilon = 0.1$, then $ATT_t^Z = 0.050$, whereas if $\epsilon = 0.7$, then $ATT_t^Z = -0.052$. This shows that whatever the sign of τ_t , the ATT^Z can be positive or negative, depending on the entire vector of regulation.

The expression (B.14) for ATC_{st}^Z is general; hence, the result for control firms in the multi-plant case is immediate: $ATC_t^Z > 0 \iff \tau_t > 0$.

Treatment Effect on Aggregate Emissions. The proof from section B.2 is general, (i.e., allows for $R_{fist} \in [0, 1]$). Hence, the result is immediate: $ATZ_t < 0 \iff \mu_t^Z > 0$.

B.6 Analytical Results on the True Plant-Level Treatment Effects for Multi-Plant Firms

In this section, we demonstrate that

Proposition 5. *When firms operate multiple plants,*

- i/ The average treatment effect on emissions for regulated plants could be positive or negative, whatever the sign of τ ;*
- ii/ The average treatment effect on emissions for unregulated plants owned by firms that own regulated plants could be positive or negative, whatever the sign of τ ;*
- iii/ The average treatment effect on emissions for unregulated plants owned by unregulated firms increases if and only if $\tau_t > 0$.*

The average treatment effect on emissions at plants from group $X \in \{TT, CT, CC\}$ within sector s can be rewritten as as

$$ATX_{st}^Z = \frac{1}{N_{st}^X} \sum_{i \in \Upsilon_s} \sum_{j \in \Omega_{ist}^X} \left[\ln e^{-\mu_t^Z R_{jfst}} + \ln \left(\frac{y_{jfst}}{y'_{jfst}} \right) \right]. \quad (\text{B.42})$$

Average Treatment Effect on Emissions from Regulated Plants. Using B.7, we have

$$\begin{aligned} ATTT_{st}^Z &= \frac{(-\mu_t^Z)}{N_{st}^{TT}} \sum_{i \in \Upsilon_{st}} \sum_{j \in \Omega_{ist}^{TT}} R_{jfst} + \frac{\left(\frac{\rho \tau_t}{\rho - 1} \right)}{N_{st}^{TT}} \sum_{i \in \Upsilon_{st}} \sum_{f \in \Omega_{ist}^T} \sum_{j \in \Omega_{fist}^{TT}} R_{jfst} \\ &- \frac{\rho - \nu}{(1 - \nu)\rho} \sum_{i \in \Upsilon_s} \frac{N_{ist}^{TT}}{N_{st}^{TT}} \ln \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho \tau_t}{\rho - 1} R_{kist}} - 1 \right) + 1 \right] \\ &- \ln \sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho \tau_t}{\rho - 1} R_{kist}} - 1 \right) + 1 \right]^{\frac{\nu(\rho - 1)}{(\nu - 1)\rho}}. \end{aligned} \quad (\text{B.43})$$

Defining $\overline{R}_{st}^{TT} \equiv \frac{1}{N_{st}^{TT}} \sum_{i \in \Upsilon_{st}} \sum_{f \in \Omega_{ist}^T} \sum_{j \in \Omega_{jist}^{TT}} R_{fist}$, simplifying the expression and combining terms yields

$$\begin{aligned} ATTT_{st}^Z &= \frac{\rho - \nu}{(1 - \nu)\rho} \sum_{i \in \Upsilon_s} \frac{N_{ist}^{TT}}{N_{st}^{TT}} \ln \left[\frac{e^{(-\mu_t^Z + \frac{\rho\tau_t}{\rho-1} \overline{R}_{st}^{TT})}}{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1} \right] \\ &- \ln \sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\frac{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1}{e^{(-\mu_t^Z + \frac{\rho\tau_t}{\rho-1} \overline{R}_{st}^{TT})}} \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}} \end{aligned} \quad (B.44)$$

Hence, $ATTT_{st}^Z > 0$ if, for all industries i ,

$$e^{(-\mu_t^Z + \frac{\rho\tau_t}{\rho-1} \overline{R}_{st}^{TT})} - 1 > \sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) \quad (B.45)$$

A simple example provides an illustration of the ambiguity of the sign in $ATTT_{st}^Z$. Suppose there is a single industry with two treated firms and several control firms. Each regulated firm owns one regulated plant and one unregulated plant. The ATTT is written as:

$$ATTT_t^Z = -\mu_t^Z + \left(\frac{\rho\tau_t}{\rho-1} \right) \left(\frac{R_1 + R_2}{2} \right) - \ln \left[\theta'_1 \left(e^{\frac{\rho\tau_t}{\rho-1} R_1} - 1 \right) + \theta'_2 \left(e^{\frac{\rho\tau_t}{\rho-1} R_2} - 1 \right) + 1 \right] \quad (B.46)$$

Suppose $\theta'_1 = .7$ and $\theta'_2 = 0.25$, $\frac{\rho\tau_t}{\rho-1} = 2$, and $\mu_t^Z = .05$. If the plants are such that $\alpha_{11} = .1$, $\alpha_{21} = .9$, $\alpha_{12} = \alpha_{22} = .5$, $R_{11} = R_{12} = 1$, $R_{21} = R_{22} = 0$, then $ATTT_t^Z = 0.089$, whereas if $\alpha_{11} = \alpha_{21} = .5$, then $ATTT_t^Z = -0.017$. Furthermore, suppose that θ 's and μ_t^Z are the same, but $\frac{\rho\tau_t}{\rho-1} = -1$. Assuming $\alpha_{11} = .9$, $\alpha_{21} = .1$, $\alpha_{12} = .2$, $\alpha_{22} = .8$, $R_{11} = R_{12} = 1$, and $R_{21} = R_{22} = 0$ yields $ATTT_t^Z = 0.017$, whereas if $\alpha_{11} = \alpha_{21} = .5$, then $ATTT_t^Z = -0.013$. This shows that whatever the sign of τ_t , the ATTT can be positive or negative, depending on the full vector of regulation.

Average Treatment Effect on Emissions from Unregulated Plants owned by Firms that own Regulated Plants. Defining $\overline{R}_{st}^{CT} \equiv \frac{1}{N_{st}^{CT}} \sum_{i \in \Upsilon_{st}} \sum_{f \in \Omega_{ist}^T} \sum_{j \in \Omega_{jist}^{CT}} R_{fist}$,

we have

$$\begin{aligned} ATCT_{st}^Z &= \frac{\rho - \nu}{(1 - \nu)\rho} \sum_{i \in \Upsilon_s} \frac{N_{ist}^{CT}}{N_{st}^{CT}} \ln \left[\frac{e^{(\frac{\rho\tau_t}{\rho-1} \overline{R}_{st}^{CT})}}{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1} \right] \\ &- \ln \sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\frac{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1}{e^{(\frac{\rho\tau_t}{\rho-1} \overline{R}_{st}^{CT})}} \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}} \end{aligned} \quad (B.47)$$

Hence, $ATCT_{st}^Z > 0$ if, for all industries i ,

$$e^{\left(\frac{\rho\tau_t}{\rho-1}\overline{R}_{st}^{-CT}\right)} - 1 > \sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1}R_{kist}} - 1 \right) \quad (\text{B.48})$$

This is essentially the same condition as the condition (B.39) for ATT_{st}^y in the multi-plant case, except for the definition of \overline{R}_{st} . The same example from Appendix B.5 indicates that, whatever the sign of τ_t , the ATCT can be positive or negative, depending on the full vector of regulation.

Average Treatment Effect on Emissions from Plants owned by Untreated Firms. Using B.7, we have

$$\begin{aligned} ATCC_{st}^Z &= - \left\{ \frac{\rho - \nu}{(1 - \nu)\rho} \sum_{i \in \Upsilon_s} \frac{N_{ist}^{CC}}{N_{st}^{CC}} \ln \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1}R_{kist}} - 1 \right) + 1 \right] \right. \\ &\quad \left. + \ln \sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1}R_{kist}} - 1 \right) + 1 \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}} \right\}. \end{aligned} \quad (\text{B.49})$$

If $\tau_t < 0$ (> 0), then the right hand side of (B.49) is negative (positive). Hence, $ATCC_{st}^Z < 0 \iff \tau_t < 0$.

B.7 Analytical Results on Bias in DD Estimator for Multi-Plant Firms

Average Treatment Effect on Firm-Level Outcomes for Regulated Firms. The DD estimate of ATT_t^v for $v \in \{y, X\}$ in the multi-plant case can be computed as $\widehat{ATT}_{st}^{v,DD} = \frac{1}{N_{st}^T} \sum_{f \in \Omega_{st}^T} \widehat{\beta}_{0,t}^{v,DD} R_{fist} = \widehat{\beta}_{0,t}^{v,DD} \overline{R}_{st}$, where $\widehat{\beta}_{0,t}^{v,DD}$ is the coefficient resulting from estimating (27) via OLS, and $\overline{R}_{st} \equiv \frac{1}{N_{st}^T} \sum_{f \in \Omega_{st}^T} R_{fist}$.

Using (B.38), we can express the bias for a given sector s as

$$\begin{aligned} Bias_{st}^{v,DD} &= \frac{\rho - \nu}{(1 - \nu)\rho} \sum_{i \in \Upsilon_s} \frac{N_{ist}^T}{N_{st}^T} \ln \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1}R_{kist}} - 1 \right) + 1 \right] \\ &\quad + \ln \sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1}R_{kist}} - 1 \right) + 1 \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}}, \end{aligned} \quad (\text{B.50})$$

which is positive if and only if $\tau_t < 0$.

Proposition 4 established that whatever the sign of τ_t , ATT_{st}^v can be positive or negative. (B.50) implies that, when $\tau_t < 0$, the DD estimator is biased up, which translates into a

bias away from zero if $ATT_{st}^v > 0$ and a bias towards zero if $ATT_{st}^v < 0$. By contrast, when $\tau_t > 0$, the DD estimator is biased down, which translates into a bias away from zero if $ATT_{st}^v < 0$ and a bias towards zero if $ATT_{st}^v > 0$. The estimated effects may even be of the wrong sign. Aggregating over sectors implies that the DD estimate of ATT_t^v can be biased towards zero or away from zero, and can even yield the opposite sign from ATT_t^v , whatever the sign of τ_t .

The DD estimator imposes that $\widehat{ATC}_{st}^{v,DD} = 0$ by assumption. Hence, the magnitude of the bias in this case equals to the true effect.

Average Treatment Effect on Aggregate Emissions. The proof in Appendix B.3 is general for $R \in [0, 1]$, so it immediately applies in the multi-plant case.

Average Treatment Effect on Plant-Level Emissions. The DD estimate of $ATTT_t^Z$ can be computed as $\widehat{ATTT}_{st}^{Z,DD} = \widehat{\beta}_{0,t}^{Z,DD,TT}$ where $\widehat{\beta}_{0,t}^{Z,DD,TT}$ is the regression coefficient resulting from estimating (28) for TT plants via OLS. We decompose the bias in the $ATTT_{st}^Z$ into two components:

$$Bias_{st}^{Z,DD,TT} = \text{Reg-Bias}_{st}^{Z,DD,TT} + \text{Spillover-Bias}_{st}^{Z,DD,TT},$$

with

$$\text{Reg-Bias}_{st}^{Z,DD,TT} = E \left[\widehat{\beta}_{0,t}^{Z,DD,TT} \right] - \left(-\mu_t^Z + \frac{\rho\tau_t}{\rho-1} \overline{R}_{st}^{TT} \right) \quad (\text{B.51})$$

and

$$\begin{aligned} \text{Spillover-Bias}_{st}^{Z,DD,TT} = & - \left\{ \frac{\rho - \nu}{(1 - \nu)\rho} \sum_{i \in \Upsilon_s} \frac{N_{ist}^{TT}}{N_{st}^{TT}} \ln \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1 \right] \right. \\ & \left. + \ln \sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1 \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}} \right\} \quad (\text{B.52}) \end{aligned}$$

We have $\text{Spillover-Bias}_{st}^{Z,DD,TT}$ positive if and only if $\tau_t > 0$, as it just encompasses spillovers from changes in the price index. The term $\text{Reg-Bias}_{st}^{Z,DD,TT}$ represents the bias in the estimate of the direct effect of the regulation on plant-level emissions of treated plants. When $\tau \neq 0$, $\text{Reg-Bias}_{st}^{Z,DD,TT}$ cannot be signed. Hence, $Bias_{st}^{Z,DD,TT}$ cannot be signed, and aggregating over sectors, $Bias_t^{Z,DD,TT}$ cannot be signed either.

The DD estimate of $ATCT_t^Z$ can be computed as $\widehat{ATCT}_{st}^{Z,DD} = \widehat{\beta}_{0,t}^{Z,DD,CT}$, where $\widehat{\beta}_{0,t}^{Z,DD,CT}$ is the regression coefficient resulting from estimating (28) for CT plants via OLS. Similarly, we decompose the bias into two components:

$$Bias_{st}^{Z,DD,CT} = \text{Reg-Bias}_{st}^{Z,DD,CT} + \text{Spillover-Bias}_{st}^{Z,DD,CT},$$

with

$$\text{Reg-Bias}_{st}^{Z,DD,CT} = E \left[\widehat{\beta_{0,t}^{Z,DD,CT}} \right] - \left(\frac{\rho\tau_t}{\rho-1} \overline{R_{st}^{CT}} \right) \quad (\text{B.53})$$

and

$$\begin{aligned} \text{Spillover-Bias}_{st}^{Z,DD,CT} &= - \left\{ \frac{\rho-\nu}{(1-\nu)\rho} \sum_{i \in \Upsilon_s} \frac{N_{ist}^{CT}}{N_{st}^{CT}} \ln \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1 \right] \right. \\ &\quad \left. + \ln \sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\frac{\rho\tau_t}{\rho-1} R_{kist}} - 1 \right) + 1 \right]^{\frac{\nu(\rho-1)}{(\nu-1)\rho}} \right\} \quad (\text{B.54}) \end{aligned}$$

Given its similarities with $\text{SpilloverBias}_{st}^{Z,DD,TT}$, $\text{SpilloverBias}_{st}^{Z,DD,CT}$ is also positive if and only if $\tau_t > 0$. When $\tau \neq 0$, then $\text{RegBias}_{st}^{Z,DD,CT}$ cannot be signed. Hence, $\text{Bias}_{st}^{Z,DD,CT}$ cannot be signed, and aggregating over sectors, $\text{Bias}_t^{Z,DD,CT}$ cannot be signed either.

B.8 Analytical Results on Bias in LA Estimator for Multi-Plant Firms

Average Treatment Effect on Firm-Level Outcomes. If condition B.29 holds, ensuring that the approximation on which the LA estimator rests is valid, then $E \left[\widehat{\beta_{0,t}^{Z,LA}} \right] = -\mu_t^Z + \frac{\rho\tau_t}{\rho-1}$, $E \left[\widehat{\beta_{2,t}^{y,LA}} \right] = \frac{(\nu-\rho)\tau_t}{(\nu-1)(1-\rho)}$, thereby $E \left[\widehat{ATT_t^{v,LA}} \right] = ATT_t^v$ and $E \left[\widehat{ATZ_t^{LA}} \right] = ATZ_t$.

Conversely, if condition B.29 does not hold, Appendix B.4 shows that the bias in the estimate of average treatment effects on outcome $v \in \{y, Z\}$ for firms in group $X \in \{T, C\}$ in sector s for the LA model can be computed as in (B.31), (B.32), (B.33), using estimates from the multi-plant version of (29) via OLS. The α s play no role in the estimated treatment effects nor in the bias, as they drop out in the computation of treatment effects. Hence, the bias in the LA model is given by the same expressions as in Appendix B.4, which indicates that the bias in the LA model cannot be signed analytically.

Average Treatment Effect on Plant-Level Emissions. The LA estimator at the plant level implies

$$\widehat{ATT_{st}^{LA}} = \widehat{\beta_t^{Z,LA}} + \widehat{\beta_{0,t}^{y,LA}} \overline{R_{st}^{TT}} + \widehat{\beta_{2,t}^{y,LA}} \sum_{i \in \Upsilon_s} \frac{N_{ist}^{TT}}{N_{st}^{TT}} \zeta_{ist_0} - \left(\widehat{\beta_{0,t}^{y,LA}} + \widehat{\beta_{2,t}^{y,LA}} \right) \zeta_{st_0},$$

where $\widehat{\beta_{0,t}^{y,LA}}$ and $\widehat{\beta_{2,t}^{y,LA}}$ are estimated from regression model (29), $\widehat{\beta_t^{Z,LA}}$ is estimated from regression model (25), $\zeta_{ist_0} \equiv \sum_{k \in \Omega_{ist}} \theta_{kist_0} R_{kist_0}$, and $\zeta_{st_0} \equiv \sum_{i \in \Upsilon_s} \Theta_{ist_0} \zeta_{ist_0}$.

The bias can be decomposed as follows

$$Bias_{st}^{ATTT,LA} = \text{Reg-Bias}_{st}^{ATTT,LA} + \text{Approx-Bias}_{st}^{ATTT,LA} \quad (\text{B.55})$$

with

$$\begin{aligned} \text{Reg-Bias}_{st}^{ATTT,LA} &= \left(E \left[\widehat{\beta}_t^{Z,LA} \right] + \mu_t^Z \right) + \left(E \left[\widehat{\beta}_{0,t}^{y,LA} \right] - \beta_{0,t}^y \right) \overline{R}_{st}^{TT} \\ &+ \left(E \left[\widehat{\beta}_{2,t}^{y,LA} \right] - \beta_{2,t}^y \right) \sum_{i \in \Upsilon_s} \frac{N_{ist}^{TT}}{N_{st}^{TT}} \zeta_{ist_0} \\ &- \left(E \left[\widehat{\beta}_{0,t}^{y,LA} \right] + E \left[\widehat{\beta}_{2,t}^{y,LA} \right] - (\beta_{0,t}^y + \beta_{2,t}^y) \right) \zeta_{st_0} \end{aligned} \quad (\text{B.56})$$

and

$$\begin{aligned} \text{Approx-Bias}_{st}^{ATTT,LA} &= \frac{\beta_{2,t}^y}{\beta_{0,t}^y} \sum_{i \in \Upsilon_s} \frac{N_{ist}^{TT}}{N_{st}^{TT}} \ln \left[\frac{e^{\beta_{0,t}^y \zeta_{ist_0}}}{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\beta_{0,t}^y R_{kist}} - 1 \right) + 1} \right] \\ &- \ln \left[\sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\frac{e^{\beta_{0,t}^y \zeta_{st_0}}}{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\beta_{0,t}^y R_{kist}} - 1 \right) + 1} \right]^{\left(\frac{\beta_{2,t}^y + \beta_{0,t}^y}{\beta_{0,t}^y} \right)} \right] \end{aligned} \quad (\text{B.57})$$

The term $\text{Reg-Bias}_{st}^{ATTT,LA}$ represents the bias in the estimated treatment effect stemming from omitted variable bias in the estimates of $\widehat{\beta}_{0,t}^{y,LA}$ and $\widehat{\beta}_{2,t}^{y,LA}$ coming from leaving Sato-Vartia-weighted productivity growth at the industry level in the error term. The term $\text{Approx-Bias}_{st}^{ATTT,LA}$ represents the bias resulting from using a linear approximation to the CES price index. In general, the sign of neither of these terms, nor the sum of them can be determined analytically.

The average treatment effect on unregulated plants owned by firms that also operate regulated plants can be computed as

$$\widehat{ATCT}_{st}^{LA} = \widehat{\beta}_{0,t}^{y,LA} \overline{R}_{st}^{CT} + \widehat{\beta}_{2,t}^{y,LA} \sum_{i \in \Upsilon_s} \frac{N_{ist}^{CT}}{N_{st}^{CT}} \zeta_{ist_0} - \left(\widehat{\beta}_{0,t}^{y,LA} + \widehat{\beta}_{2,t}^{y,LA} \right) \zeta_{st_0}.$$

The bias can be decomposed as follows

$$Bias_{st}^{ATCT,LA} = \text{Reg-Bias}_{st}^{ATCT,LA} + \text{Approx-Bias}_{st}^{ATCT,LA} \quad (\text{B.58})$$

with

$$\begin{aligned} \text{Reg-Bias}_{st}^{ATCT,LA} &= \left(E \left[\widehat{\beta_{0,t}^{y,LA}} \right] - \beta_{0,t}^y \right) \overline{R}_{st}^{CT} + \left(E \left[\widehat{\beta_{2,t}^{y,LA}} \right] - \beta_{2,t}^y \right) \sum_{i \in \Upsilon_s} \frac{N_{ist}^{CT}}{N_{st}^{CT}} \zeta_{ist_0} \\ &- \left(E \left[\widehat{\beta_{0,t}^{y,LA}} \right] + E \left[\widehat{\beta_{2,t}^{y,LA}} \right] - (\beta_{0,t}^y + \beta_{2,t}^y) \right) \zeta_{st_0} \end{aligned} \quad (\text{B.59})$$

and

$$\begin{aligned} \text{Approx-Bias}_{st}^{ATCT,LA} &= \frac{\beta_{2,t}^y}{\beta_{0,t}^y} \sum_{i \in \Upsilon_s} \frac{N_{ist}^{CT}}{N_{st}^{CT}} \ln \left[\frac{e^{\beta_{0,t}^y \zeta_{ist_0}}}{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\beta_{0,t}^y R_{kist}} - 1 \right) + 1} \right] \\ &- \ln \left[\sum_{i \in \Upsilon_s} \theta'_{ist} \left[\frac{e^{\beta_{0,t}^y \zeta_{st_0}}}{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} \left(e^{\beta_{0,t}^y R_{kist}} - 1 \right) + 1} \right]^{\left(\frac{\beta_{2,t}^y + \beta_{0,t}^y}{\beta_{0,t}^y} \right)} \right] \end{aligned} \quad (\text{B.60})$$

The term $\text{Reg-Bias}_{st}^{ATCT,LA}$ represents the bias in the estimated treatment effect stemming from omitted variable bias in the estimates of $\widehat{\beta_{0,t}^{y,LA}}$ and $\widehat{\beta_{2,t}^{y,LA}}$, as above. The term $\text{Approx-Bias}_{st}^{ATCT,LA}$ represents the bias resulting from using a linear approximation to the CES price index. In general, the sign of neither of these terms, nor the sum of them can be determined analytically.

The average treatment effect on emissions of plants owned by control firms can be computed as

$$\widehat{ATCC}_{st}^{LA} = \widehat{\beta_{2,t}^{y,LA}} \sum_{i \in \Upsilon_s} \frac{N_{ist}^{CC}}{N_{st}^{CC}} \zeta_{ist_0} - \left(\widehat{\beta_{0,t}^{y,LA}} + \widehat{\beta_{2,t}^{y,LA}} \right) \zeta_{st_0}.$$

The bias can be decomposed as follows

$$\text{Bias}_{st}^{ATCC,LA} = \text{Reg-Bias}_{st}^{ATCC,LA} + \text{Approx-Bias}_{st}^{ATCC,LA} \quad (\text{B.61})$$

with

$$\begin{aligned} \text{Reg-Bias}_{st}^{ATCC,LA} &= \left(E \left[\widehat{\beta_{2,t}^{y,LA}} \right] - \beta_{2,t}^y \right) \sum_{i \in \Upsilon_s} \frac{N_{ist}^{CC}}{N_{st}^{CC}} \zeta_{ist_0} \\ &- \left(E \left[\widehat{\beta_{0,t}^{y,LA}} \right] + E \left[\widehat{\beta_{2,t}^{y,LA}} \right] - (\beta_{0,t}^y + \beta_{2,t}^y) \right) \zeta_{st_0} \end{aligned} \quad (\text{B.62})$$

and

$$\begin{aligned} \text{Approx-Bias}_{st}^{ATCC,LA} &= \frac{\beta_{2,t}^y}{\beta_{0,t}^y} \sum_{i \in \Upsilon_s} \frac{N_{ist}^{CC}}{N_{st}^{CC}} \ln \left[\frac{e^{\beta_{0,t}^y \zeta_{ist_0}}}{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} (e^{\beta_{0,t}^y R_{kist}} - 1) + 1} \right] \\ &- \ln \left[\sum_{i \in \Upsilon_s} \Theta'_{ist} \left[\frac{e^{\beta_{0,t}^y \zeta_{st_0}}}{\sum_{k \in \Omega_{ist}^T} \theta'_{kist} (e^{\beta_{0,t}^y R_{kist}} - 1) + 1} \right]^{\left(\frac{\beta_{2,t}^y + \beta_{0,t}^y}{\beta_{0,t}^y} \right)} \right] \quad (\text{B.63}) \end{aligned}$$

The term $\text{Reg-Bias}_{st}^{ATCC,LA}$ represents the bias in the estimated treatment effect stemming from omitted variable bias in the estimates of $\widehat{\beta_{0,t}^{y,LA}}$ and $\widehat{\beta_{2,t}^{y,LA}}$, as above. The term $\text{Approx-Bias}_{st}^{ATCC,LA}$ represents the bias resulting from using a linear approximation to the CES price index. In general, the sign of neither of these terms, nor the sum of them can be determined analytically.


C Monte Carlo Experiments

In this appendix, we present additional results from Monte Carlo experiments not reported in the main text.

C.1 Monte Carlo Experiments for Single-Plant Firms

In the paper, we show results from Monte-Carlo experiments assuming that firms can own multiple plants. Appendix B reveals that if the data generating process assumes that all firms own a single plant, we can sign analytically the biases at least for the DD estimator. By contrast, if market shares do not remain at their pre-regulation values, the sign of the LA estimator biases is indeterminate. Figure C.7 shows the biases in estimating average treatment effects from each estimator assuming that all firms own a single plant and follow a random walk productivity growth path.

Figure C.7: Treatment Effects for revenues and emissions at the firm level and in aggregate, Single-Plant Firms



Notes: Subfigures in panel a) plot average estimates of treatment effects on revenues (y) and emissions (z) as a percentage of observed outcomes for each estimator across 100 replications against the average of the true metric, for each parameter combination. Left panels show the ATTs for each parameter combination, whereas right panels show the ATCs, all at the firm level. Panel b) plots the average aggregate effect on total emissions relative to the counterfactual unregulated equilibrium. The black line corresponds to the 45-degree line. Data generating process assumes single-plant firms and random walk productivity growth.

C.2 Monte-Carlo Experiments for Oligopoly

In the model, we assume that firms do not internalize the effect of their pricing on the industry-wide price index. However, with a discrete number of firms, some of whom command nontrivial market shares, it could be that, in reality, at least some firms do internalize the effect of their pricing behavior on the CES price index. In this case, our estimator would be misspecified. To explore the potential role of misspecification bias with respect to strategic behavior, we run an adversarial test on our empirical strategy by simulating the data using Bertrand-Nash pricing behavior, and leaving all other features of the simulations the same.

To simulate the economy assuming firms engage in Bertrand-Nash pricing, we consider the following pricing rule since firms internalize the effect of their pricing decision on the CES price index:

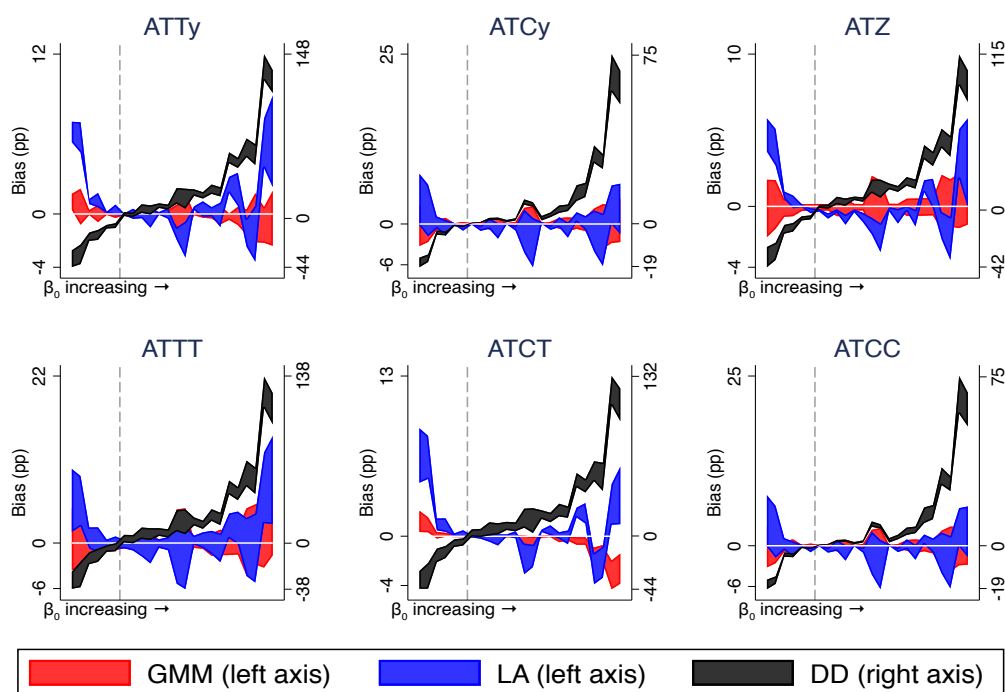
$$p_{fit}^O = c_{fit} \left[\frac{(\nu - 1) - (\nu - \rho) \theta_{fit} + (1 - \rho) \nu \theta_{fit} \Theta_{it}}{\rho(\nu - 1) - (\nu - \rho) \theta_{fit} + (1 - \rho) \nu \theta_{fit} \Theta_{it}} \right], \quad (\text{C.64})$$

where $\theta_{fit} = (p_{fit}/P_{it})^{\frac{-\rho}{1-\rho}}$ describes the response of the within-industry price index to firm f 's price change, and $\Theta_{it} = (P_{it}/\Psi_t)^{\frac{-\nu}{1-\nu}}$ describes the response of the sector price index to industry i 's price index change. The pricing rule thus varies with firm f 's market share in industry i and industry i 's market share in sector s at time t .

We use the same empirical strategies as presented in the paper for multi-plant firms – our GMM estimator, the DD model, and the LA model – and compare their estimated ATTs with the true ATTs. Figure C.8 shows the interquartile distribution of estimation errors by estimator and parameter combination assuming oligopolistic behaviors. We find that, even if firms play Bertrand-Nash, our procedure still delivers estimates of treatment effects that matches the true values quite well. Indeed, the interquartile of the distribution of errors for our GMM estimator always covers zero, except for the ATCT for highly negative τ_t . In any case, extremes of the interquartile distribution are never more than a few percentage points away from zero. By contrast, the LA estimator performs less well since its interquartile distribution of estimation errors often does not recover zero (see for instance, subfigures representing ATT^y , ATZ, ATTT, and ATCT) and is further away from zero for some parameter combinations. The DD estimator appears strongly biased: estimation biases are measured on the right axis of each subfigure, implying large magnitudes. As predicted, the sign of the bias is opposite to the sign of τ_t . These patterns are similar to the ones observed under monopolistic competition.

The good performance of our GMM estimator is reminiscent of the findings in Head & Mayer (2022), where it was found that even if firms play Bertrand-Nash, estimation based on the assumption of monopolistic competition performs incredibly well.

Figure C.8: Bias in Average Treatment Effects in Monte Carlo Simulations, Oligopoly



Notes: Subfigures plot the interquartile distribution of estimation errors by estimator and parameter combination for average and aggregate treatment effects across 100 replications for a DGP that assumes oligopoly. On the x-axis, parameter combinations are ordered by $\beta_0 = \frac{\rho\tau_t}{\rho-1}$. GMM and LA estimation biases are measured on the left axis, while DD estimation biases are measured on the right axis. In all subfigures, parameter combinations to the left (right) of the vertical dashed line indicate simulations for which $\tau_t > 0$ (< 0).

D Data Appendix

D.1 FICUS-FARE

Balance sheet information for the universe of French firms was retrieved from the FICUS (Annual structural statistics of companies from the SUSE scheme, 1994-2007) and FARE (Annual structural statistics of companies from the ESANE scheme, 2008-2016). These firm-level data originate from a fiscal source, since firms need to declare their profits to the tax authorities. Firm-level data from FICUS/FARE were linked to other data sets based on the unique identifier of French firms (SIREN).

We identify industries and sectors based on the activity code APE (*Activité Principale de l'Établissement*), from the national activity nomenclature, namely the *Nomenclature des Activités Françaises* (NAF). These 4-digit codes correspond to an industry, whereas the first two digits of the codes are common within a sector. The NAF classification was revised in 2003 (from the “NAF93” to the “NAFRev.1”) and again in 2008 (to the “NAFRev.2”). If the conversion from the NAF93 to the NAFRev1 is straightforward, the revision in 2008 deeply modified the within-sector industry decomposition and resulted in a many-to-many mapping between NAFRev.1 and NAFRev.2. Since our analysis rests on a structure of industries and sectors, we must assign one activity code to each firm over time. We use the NAFRev.1 revisions in most of the analysis. If a firm is active both before and after 2008, it is assigned two stable activity codes, one for each revision. Indeed, we assign to firms that switch industry codes their modal code within a revision. If a firm is not active in one of the period, it is either assigned a 1-to-1 match in nomenclatures (if available) or a code in the NAF classification that is observed most frequently for firms with the same industry code.

We include all firms in all industries with 2-digit codes between 15 - 37 in the NAF Rev.1. We exclude extractive industries (codes 10 to 14) and energy production and distribution (40 and 41). We group aggregates of 2-digit APE codes into “sectors”, as in Harrigan et al. (2018b). We consider the stable firm-level 4-digit APE code as the industry.

D.2 EACEI

The information on energy use comes from the EACEI (*Enquête sur les consommations d'énergie dans l'industrie*) and EACEI-IAA (for agro-industry) surveys. These are surveys of manufacturing establishments (identified by a unique identifier called SIRET, whose first 9 digits identify the SIREN of the firm) that provide information on the consumption (quantity and value) of energy, broken down by energy types: electricity (bought and self-generated), steam, natural gas, other types of gas, coal, lignite, coke, propane/butane, heavy fuel oil, heating oil and other petroleum products. All establishments with more than 250 employees receive the survey each year, whereas only a sample of establishments with 20 or more employees (stratified by industries, number of employees, and region) receive it. The response rate is nearly 90% (see Marin & Vona 2021). Figure D.9 shows that plants covered by the EACEI represent only 10% of the number of manufacturing plants

in France, but 50% of the labor employed in manufacturing and 80% of manufacturing emissions (using sector-year averages of emission intensity to impute missing data).

We compute CO2 emissions at the plant level by combining the quantity of energy consumed with an energy-specific conversion factor that indicates the amount of CO2 released in the atmosphere when the type of energy considered is used. We thus focus on energy-related CO2 emissions released by each plant during production through the combustion of fossil energy or through the indirect emissions related to electricity bought from the grid or steam bought externally. We thus ignore any emissions released in the upstream stages of extraction, transformation and transportation of energy. We use CO2 emission factors for the different energy types from ADEME’s “Base Carbone”.⁴⁶ Whereas the conversion factors are constant over time for most energy types, given the proportionality between CO2 emissions and the carbon content of fossil fuels, the amount of CO2 emissions associated with the use of electricity from the grid is time varying and depends on the energy mix of the electricity sector in France. Table D.2 reports these conversion factors.

Table D.2: CO2 Conversion Factors

Fuel name	Unit	Conversion Factor
Coal	t	3.07
Lignite	t	1.72
Coke of coal	t	3.03
Petroleum coke	t	3.1
Gas from the grid	MWh	0.169
Non-natural Gas	MWh	0.469
Butane/Propane	t	2.965
Heavy fuel oil	t	3.14
Heating oil	L	0.00268
Steam	t	0.113
Electricity		
2012	MWh	0.0552
2013	MWh	0.0526
2014	MWh	0.0492
2015	MWh	0.0408
2016	MWh	0.0379


Notes: These conversion factors have been computed by the authors using ADEME’s “Base Carbone”. The units for fuels are in tons (t), megawatt hours (MWh) or liters (L). The conversion is toward tons of CO2.

To clean the data, we first identify whether there are within-plant observations that are 50 times smaller or larger than others over time. If so, we drop the entire history of plant

⁴⁶We use CO2 conversion factors from the document published by ADEME entitled “Base Carbone: Documentation des Facteurs d’Emissions de la Base Carbone” in its version 11.2.0, available from March 2015.

CO2 emissions. Next, we identify within-plant year-to-year changes that are equivalent to being multiplied by 10 or divided by 10. If so, we again drop the entire plant-level history.

Figure D.9: Share of Plants Covered by the EACEI Survey



Notes: Figure plots the shares of plants, labor share, and emissions shares of plants included in the EACEI survey by year. Emissions share is computed by imputing emissions for plants excluded from the EACEI survey.

D.3 EUTL

To assess the treatment status of each firm and plant, we use the European Union Transaction Log (EUTL), which is the central reporting tool for the EU ETS. Transactions are reported through the EUTL with a delay of three years. The data can be downloaded from the EUTL webpage (<https://ec.europa.eu/clima/ets/>). We downloaded two registries, one for 2014 and the other for 2018.

Within the EUTL, an installation is a regulated entity. It faces the obligation to surrender allowances at least equal to its verified emissions of the previous year to the regulatory authority. Installations either receive these allowances for free or buy them on the allowances market. To be able to receive, transfer, and surrender allowances, each installation must be represented by an operator holding account (OHA). For each account, the EUTL provides a primary contact for the account holder with an address. Each installation receives a unique identifier and a registry, which corresponds to its country of location. In some cases, the EUTL will provide information on companies related to the installation, that is on the SIREN in the French context. SIREN identifiers are usually reported for manufacturing industries, but not for the power sector and for heating units. To illustrate, from the registry of 2014, we obtain 1256 account holders with SIREN information and

Table D.3: Treatment by Sector

	# Firms	# Plants
<i>Manufacturing</i>		
Motor vehicles & other transport equipment	11	13
Chemicals & Pharmaceuticals	85	126
Computer, Electronic & Optical products	4	4
Electrical equipment	2	2
Food, Beverages & Tobacco	104	148
Machinery & Equipment	7	9
Basic & Fabricated Metal products	40	57
Rubber, Plastic & Non-metallic mineral products	88	198
Textiles & Apparel	17	17
Wood & Paper products	99	116
Other manufacturing	8	9
Subtotal	465	699
<i>Non-Manufacturing/Unknown</i>	381	565
Total	846	1264

130 without. To complement the original file, we use the Ownership links and enhanced EUTL dataset from Jaraite, Jong, Kazukauskas, Zaklan and Zeitlberger (2016).⁴⁷ From the registry of 2018, we obtain 995 account holders with SIREN and 320 without, some of them being identical to the ones from the registry of 2014.


To assign a treatment status to each plant, we must match the EUTL installations with the list of SIRET owned by a firm (SIREN) in all manufacturing sectors. We first retrieve the list of plants and their locations (city code, address) from the “Stock of Establishments” dataset from the French statistical agency INSEE. The INSEE agency produces this dataset and the information on plant-level labor force using a combination of CLAP (knowledge on local production entities), DADS (employer-employee dataset), and EPURE (data from the French social security system). We then merge the installations from EUTL with this list using SIREN and city codes. Next, for the installations that match with several SIRET, we use the street names and a web scraping python code to select among these different SIRET.

The number of regulated plants and firms by sector (using the nomenclature NAF Rev.1) are reported in Table D.3. We count 465 firms and 699 plants ever regulated in the manufacturing sectors.

Figure D.10 shows the distribution of plant-level CO₂ emissions in 2004, as well as the share of plants being regulated within each bin of CO₂ emissions level in 2004. We note that the regulation share never reaches 100% even for the most polluting plants.

⁴⁷The dataset can be downloaded from <https://cadmus.eui.eu/handle/1814/64596>.


Figure D.10: Regulation status by plant-level CO₂ emissions in 2004



D.4 Imports and Exports

We use the Customs data for the export sales by firm-product-destination for each year of the sample, and aggregate at the firm level. We use BACI from CEPII for aggregate trade flows and select all flows toward France. Figure D.11 shows that roughly a third of manufacturing sales in France comes from abroad.

Figure D.11: Sales on French Market by Firm Origin



Notes: This figure plots the total sales in France across all manufacturing sectors for French firms vs foreign firms. Foreign sales are computed from BACI.

D.5 Variable Definitions

Revenue is total sales, including exports. In FICUS, this is CATOTAL, whereas in FARE this is REDI_R310. Domestic sales are the sales realized on the domestic market (CAFRANC in FICUS, REDI_R420 in FARE). We compute them as the difference between total sales and export sales from the French Customs dataset.

Employment is the full-time equivalent of the number of directly employed workers by the plant averaged over the year. We use the variable reported in the Stock d’Etablissements.

Energy consumption per energy type is defined as the quantity of energy used by a plant. In EACEI, the variable is CONS_UP for the surveys from 1994 to 1999, and CSUP for the surveys from 2000 onward.

D.6 Imputing Emissions

Since we do not observe all plants in the EACEI surveys, we can impute emissions from missing plants using average sector-year emission intensity. These imputations are not used to estimate reduced-form parameters, but only to build counterfactual emissions.

Given our theoretical model, emissions at the plant-level are given by

$$z_{jfst} = \underbrace{\frac{\rho\gamma\kappa_t}{w_t^E}}_{\equiv x_t} e^{\widehat{\beta}_{10}R_{jfst}} \alpha_{jfst} y_{jfst}$$

where we have used our estimate $\widehat{\beta}_{10}$ in place of the underlying structural parameter $-\mu_z$. Note that the emission intensity term $x_t = \frac{\rho\gamma\kappa_t}{w_t^E}$ does not vary by firm within a year. Hence, we can estimate x_t by taking averages over the year, by sector,

$$\widehat{x}_{st} = \frac{1}{N_{st}} \sum_j \frac{z_{jfst}}{\alpha_{jfst} * y_{jfst}} e^{-\widehat{\beta}_{10}R_{jfst}}$$

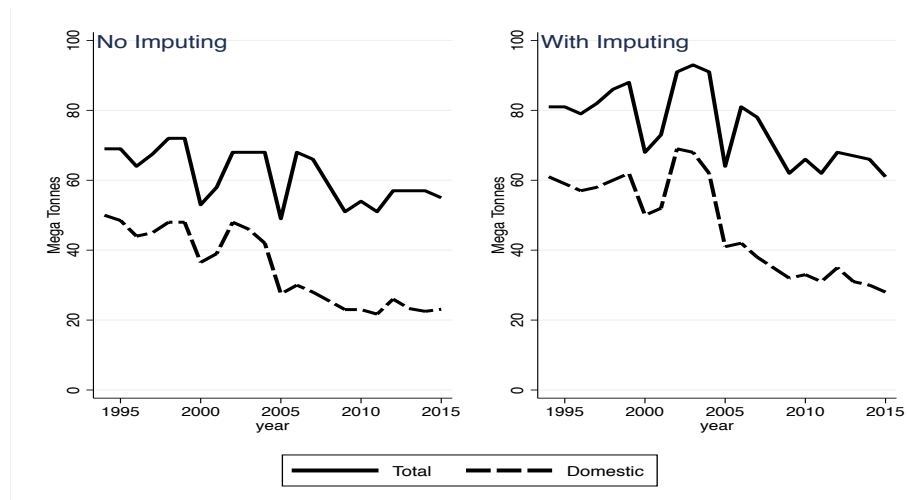
Then for any plant whose emissions we don’t observe, we can compute

$$\widehat{z}_{jfst} = \widehat{x}_{st} e^{\widehat{\beta}_{10}R_{jfst}} \alpha_{jfst} y_{jfst}.$$

Additionally, we compute emissions related to production for the domestic market only, by decomposing total sales into domestic and export sales: $y_{jfst} \equiv y_{jfst}^{dom} + y_{jfst}^{for}$, and by attributing emissions to domestic sales:

$$\widehat{z}_{jfst}^{dom} = \frac{z_{jfst}}{y_{jfst}} y_{jfst}^{dom}.$$

Figure D.12: Emissions, with and without Imputations

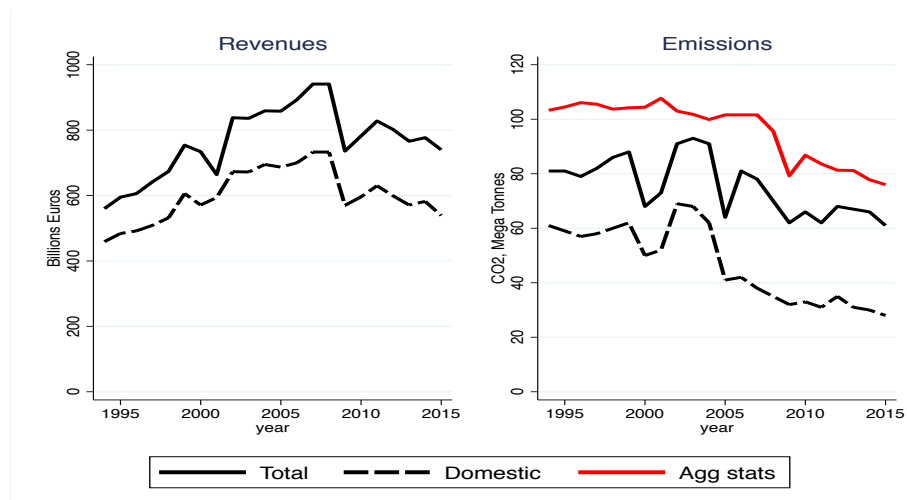


Notes: Figure plots total (solid) and domestic (dashed) emissions, without (left) and with (right) imputing emissions for plants not surveyed by the EACEI.

Figure D.12 plots total CO₂ emissions (solid) vs CO₂ generated in production for domestic sales (dashed) with imputing emissions (right) and without imputing emissions (left). The jumps in the CO₂ emissions series are related to large metalworking factories seeing their legal entities changed (resulting in a change in SIREN), which results in these plants being missing in the EACEI survey the year after.

Figure D.13 shows that aggregate CO₂ emissions from the EACEI surveys, with the imputation presented above, follows the same trend as the reported emissions from all manufacturing sectors in the National Emissions Inventory established by the French government.

Figure D.13: Total vs Domestic Sales and Emissions for French Firms



Notes: Figure plots the total revenues vs revenues on the French market (domestic) for French firms (left) and total CO₂ emissions vs CO₂ generated in production for domestic sales for French firms (right). CO₂ emissions are imputed for plants not included in the EACEI survey. The red line plots the total CO₂ emissions from all manufacturing sectors reported in the National Emissions Inventory of France.