Climate Policy, Industry Concentration and Technology Adoption: Evidence from the French Cement Industry*

Niklas Schoch

This version: March 9, 2026
Preliminary Draft - Please do not circulate

Abstract

Industry decarbonisation remains a major challenge to reduce greenhouse gas emissions. The most polluting industries are typically characterised by market concentration. This paper investigates the welfare implications of carbon pricing in concentrated industries when firms can adopt emissions-saving technologies. I show that the optimal carbon tax level must balance output suppression with sufficient incentives to adopt available green technologies. To quantify this trade-off, I develop a dynamic model of imperfect competition and endogenous production technology choice, using recent developments in the modelling of dynamic continuous investment choices. I estimate the model using data from the French cement industry, exploiting a detailed dataset on low-carbon investments and the recent escalation in the EU carbon price to estimate abatement investment costs. I use the estimated model to simulate investment and industry paths under varying levels of a carbon tax and technology cost subsidies. In this setting, the welfare-maximising carbon tax lies between the level implied by a static oligopoly analysis and the Pigouvian level. This reflects the trade-off between further output suppression (Buchanan-critique) and sufficient incentives for green investments. In addition, I show that the combination of tax and subsidy provides sufficient investment incentives without triggering a strong output contraction from excessive taxation. As a consequence, I find supplementing the statically optimal tax with a technology cost subsidy yields a modest welfare gain above what the optimal tax alone can achieve.

Keywords: Climate policy, heavy industry, market power, carbon pricing, green technology adoption

*SciencesPo Paris, Department of Economics, 28 Rue des Saints-Pères, F-75007 Paris; niklas.schoch@sciencespo.fr.

I am indebted to Johannes Boehm and Stefan Pollinger for their incredible supervision and support throughout the entire project. I am also deeply grateful to Gautam Gowrisankaran for hosting me at Columbia University, where he provided invaluable feedback and guidance, especially considering the technical aspects of this paper. Further, I would like to extend a special thank you to Isabelle Méjean, who, in addition to countless helpful comments, provided me with access to the data, without which this project would not have been possible. I would like to thank Haoge Chang, Nicolas Cœurdacier, Antoine Ferey, Junnan He, Marleen Marra, Thierry Mayer, Kevin O'Rourke, Franz Ostrizek, Noémie Pinardon-Touati, Mar Reguant, Emma Du Puy, Andrey Simonov, and seminar participants at Sciences Po, Columbia University, and IAE-CSIC Barcelona for their helpful comments. The computing resources for this project were provided by the Nuvolos Fellowship Programme 2026.

1 Introduction

Curbing industry emissions is a significant obstacle to decarbonising global economic activity. Within the EU, emissions from industrial installations accounted for more than 20% of greenhouse gas emissions (EP, 2024). To achieve decarbonisation, carbon pricing is a central element of the climate policy strategies in many countries. Its objective is to reduce carbon emissions by taxing emissions-intensive production processes. One major avenue for achieving emission reductions is upgrading technologies to lower emission intensity. The most polluting industries targeted by climate policies, for example, the Cement, Steel, and Chemical industries, are typically highly concentrated (Koltay & Lőrincz, 2021), which affects firms' upgrade decisions and, in turn, the welfare implications of climate policies (Fowlie, 2010; Hoarau & Ponssard, 2025). An efficient design of carbon pricing policies must therefore take market structure into account.

Given the unprecedented scope of global climate policy regulations in recent years and the growing backlash over concerns about deindustrialisation, understanding industry responses is critical to informing the policy debate on suitable future climate policies. To contribute to this discussion, this paper analyses the welfare effects and firm responses to climate policy in concentrated markets, explicitly including endogenous investments in abatement technologies.

To do so, I propose a dynamic model of imperfect competition and strategic abatement technology choice, and estimate it using data on production and abatement investment in the French cement industry. I model the competitive structure as a regionally segmented Cournot model, where firms compete under capacity constraints (Ryan, 2012; Fowlie et al., 2016) and face a cost for their emissions due to environmental regulations. Dynamically, firms can invest in green technologies that reduce the emission intensity of their output. Due to the price of emissions, reducing their emission intensity reduces their marginal costs, thereby altering their competitive position in future periods. To compute the investment responses and welfare implications of varying levels of a carbon tax and a technology cost subsidy, I solve for the Markov-perfect Nash equilibrium of the dynamic game and forward simulate the expected industry path.

This allows me to evaluate counterfactual outcomes under varying environmental regulations. Conventional wisdom holds that, in the presence of market power, taxation of externalities should be set below the Pigouvian level (the social cost of carbon) because market power already suppresses output below the socially optimal level. However, this holds only if a firm's sole response margin to carbon taxation is an adjustment in output levels. In reality, a major objective of carbon taxation is to incentivise the adoption of green technology, reducing the emission intensity of output. My results show that climate policy in concentrated industries must strike a balance between output suppression (Buchanan critique) and sufficient incentives to adopt low-emission technologies. As a consequence, the welfare-maximising carbon tax level lies between the Pigouvian level and the welfare-maximising tax implied by the static analysis proposed by Buchanan (1969). For a social cost of carbon of €125/ton, I estimate an optimal carbon tax of €105/ton, compared to €60/ton implied by a static oligopoly analysis. Setting the tax at the statically optimal level, ignoring the technology adoption margin, would result in a welfare loss of 5%. Furthermore, my analysis confirms the finding of Hoarau and Ponssard (2025) that firms with market power have a lower incentive to invest in green technology than the socially optimal level. Thus, there are multiple distortions present: suppressed output due to market power and environmental externalities, and additionally, a suboptimal incentive for firms to invest in

emission-reducing technologies. As a result, a policy setting only a carbon price cannot achieve the welfare optimum. By correcting the suboptimal investment incentives, I estimate that an investment cost subsidy for green technology of 55%, in combination with the statically optimal tax, yields a 1.9% welfare gain above what the optimal tax alone can achieve.

This study innovates on current structural IO models of environmental regulation by introducing an explicit endogenous abatement investment decision into a dynamic oligopoly game. To allow for a continuous investment choice, I employ recent advances in the empirical industrial organisation literature and apply them to a dynamic oligopoly. Constructing a rich dataset on abatement investment in the French cement industry under EU ETS regulation enables me to estimate the model, including the abatement investment cost function. Thus, I can quantify the equilibrium adoption of green technology in a concentrated industry within a full structural model, complementing earlier theoretical work on firms' upgrading incentives. My results can explain the robust finding of limited adverse effects of carbon pricing in impact evaluations. I show that firms' ability to invest in emission abatement limits the economic burden of carbon pricing for both firms and consumers.

The French cement industry provides an ideal setting for this study. Firstly, the high inland transport costs cause firms to compete locally, which allows a good characterisation of the competitive environment faced by each firm (Miller & Osborne, 2014; Fowlie et al., 2016).[1] Using administrative data on production from 2012-2019, I can exploit the well-defined competitive setting and identify the static production cost parameters from fluctuations in demand and output responses between regions and over time, using the estimator proposed by Ryan (2012). The production of a standardised product under capacity constraints, competing with a few large players, reflects the industry structure of many polluting industries, such as steel, chemicals, and fertilisers. Thus, I hope the insights from the cement market carry over to other sectors targeted by the same climate policies.

Secondly, the European Union has introduced a carbon price through the EU Emissions Trading System (EU ETS), under which all firms in the French cement industry fall. By compiling a novel dataset that combines administrative data on green investments with hand-collected data on recent firm announcements, I have detailed information on the compliance investment behaviour of an industry subject to a stringent carbon price. Figure 1 shows that a wave of low-carbon investments has accompanied the recent escalation of the carbon price within the EU ETS. I use the variation in firms' investment responses to carbon price changes to quantify the cost of adopting green technology for French cement firms using a GMM nested fixed-point approach.

The rest of the paper is structured as follows. I discuss the related literature in Section 2 and the institutional details of the cement industry in Section 3, before laying out the model in Section 4, data in Section 5, estimation in Section 6 and the counterfactual results (Section 7). The last section concludes.

2 Literature

I mainly relate to three strands of literature. First and foremost, I contribute to the structural IO literature on the implications of environmental regulations (among others: Ryan, 2012; Fowlie et al., 2016; Ganapati et al., 2020; Gowrisankaran et al., 2024; Toyama, 2024; Gowrisankaran et al., 2025), by introducing endogenous green technology upgrading of firms. Low-carbon technology adoption, a

[1] I account for import competition via ship, with more details discussed in the following sections.

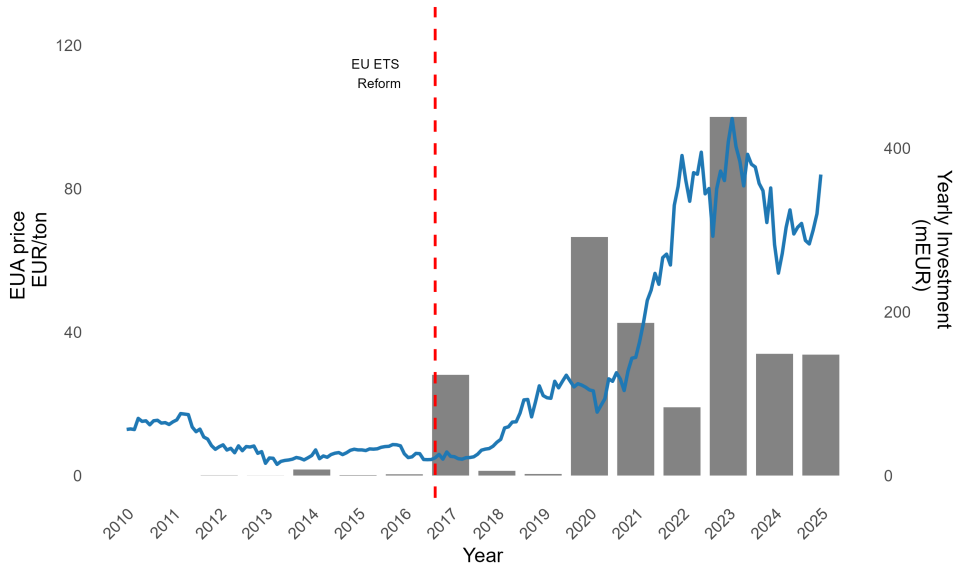


Figure 1: EU ETS price and green investment in the French cement sector

primary objective of climate policies, is a significant response margin for firms facing environmental regulations, as evidenced in impact evaluations of the EU ETS (Colmer et al., 2025). Therefore, welfare assessments must reflect this margin. However, the endogenous choice of abatement investment in a structural model has been studied only in competitive settings, such as the electricity sector (Toyama, 2024; Gowrisankaran et al., 2025). To reflect the market concentration typical of heavy industries, I advance the literature by introducing the investment margin to a dynamic oligopoly setting. In particular, I use the novel framework for continuous investment choices of Gowrisankaran and Schmidt-Dengler (2024). Thus far, it has only been applied in a single-agent setting (Gowrisankaran et al., 2024). I incorporate it into a full multiple-agent dynamic game. Hence, I explicitly allow for strategic interactions in firms' upgrading decisions. To do so, I apply the algorithm proposed by Gowrisankaran and Schmidt-Dengler (2024) to rapidly compute choice probabilities in a continuous investment setting.

Furthermore, my results complement the literature employing econometric techniques to estimate observed effects of carbon pricing (among others: Fabra & Reguant, 2014; Martin et al., 2014; Calel & Dechezleprêtre, 2016; Wagner et al., 2016; Calel, 2020; Dechezleprêtre et al., 2023; Colmer et al., 2025). Most of these empirical studies find little evidence for output contractions or other adverse effects of climate policies (among others: Wagner et al., 2016; Colmer et al., 2025). Instead, they document adjustments in energy use (Martin et al., 2014) and capital stock (Colmer et al., 2025). By embedding these adjustment mechanisms into a structural model, I provide a rationale for their findings, showing that firms' low-carbon investments can significantly reduce the economic burden of carbon pricing on firms and consumers. Additionally, this study illustrates the mechanisms and drivers behind the reduced-form results, allowing to derive targeted policy implications.

Thirdly, I relate to the early theoretical literature on technology investments and environmental regulations (among others: Delbono & Denicolo, 1991; Laffont & Tirole, 1996; Yi, 1999; Montero, 2002; Requate & Unold, 2003; Requate, 2005); see Phaneuf and Requate (2016) for a review. Most closely related is a recent study by Hoarau and Ponssard (2025), who show, in a stylised Cournot setting, that the incentives to adopt carbon-capture technology are suboptimal due to imperfect competition. The full structural model allows me to quantify the costs and the equilibrium adoption of green technol-

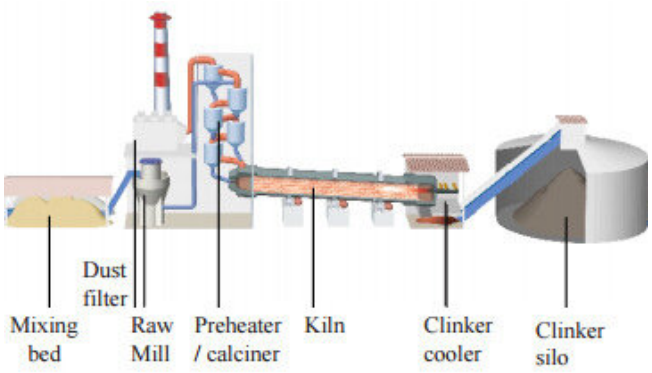


Figure 2: Cement production line (Marie et al., 2018)

ogy in a continuous choice setting, as well as the welfare implications of a wide range of alternative policies.

3 Institutional Details

This study analyses the hydraulic cement market in France. Mixed with sand, water, and other aggregates, hydraulic cement forms concrete, a material used extensively in construction for residential and commercial buildings, as well as bridges and roads. Hydraulic cement is produced by heating limestone to high temperatures. Figure 2 illustrates a typical cement production process. This involves feeding the wet limestone into the production kiln through a series of funnels designed to use exhaust heat to pre-dry the raw material. After reaching over 1400°C, limestone turns into cement clinker, which is ground up and mixed with additives such as gypsum to produce cement, a dry fine powder. Within the EU, approximately 90% of cement is sold in bulk to construction firms. In contrast, the remaining 10% is sold directly to end consumers (CEMBUREAU, 2025).[2]

Abatement Options in the Cement Industry. The cement industry is a significant emitter of greenhouse gases in the EU. In France, it accounts for roughly 15% of manufacturing GHG emissions. In the last decade, the emission intensity of cement production has fluctuated slightly around 0.6 tons of CO_2-equivalent per ton of cement produced (ADEME, 2023). The emissions in the production process arise from two sources: chemical processes that occur when limestone is heated (process emissions) and heat generation, which is typically achieved using fossil fuels (energy emissions). Roughly 60% of emissions are from chemical processes, and the remaining 40% are from energy use.

Previous studies focused on energy efficiency of new entrants to the market as a main abatement lever (Fowlie et al., 2016). However, the French cement industry almost exclusively uses dry production processes, operating rotary kilns equipped with preheater and precalciner technology, which is the most efficient commercially employed technology. Given the considerable efficiency gains, the French cement sector had already broadly adopted this technology before the introduction of the EU ETS.[3] Therefore, investments into energy efficiency are no longer a substantial abatement margin in

[2]The cement discussed here refers to grey hydraulic cement, denoted by the product code 23.51.12 (CPF rev. 2.1). This classification predominantly includes Portland cement, accounting for over 85% of production volumes within this category.

[3]The few remaining wet kilns are retrofitted within my sample period. They will be included in the estimation alongside other low-carbon investments.

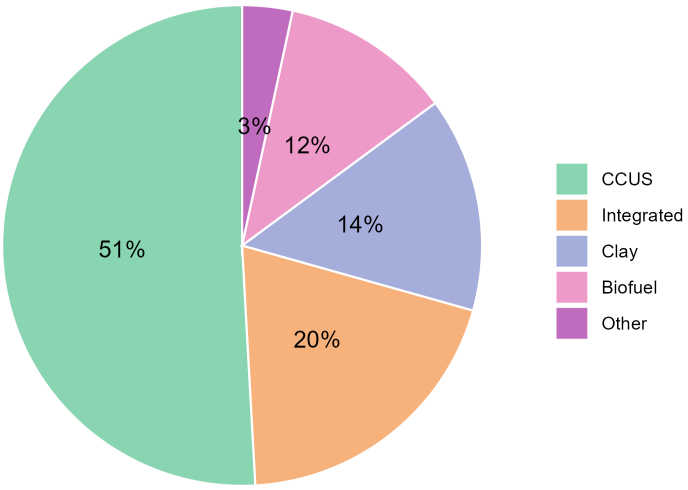


Figure 3: Share of emissions abated by technology

my study period. Instead, starting from this technology standard, I observe three abatement levers being employed to address both process and energy emissions. Firstly, the fossil fuel used to heat the kiln can be replaced by biofuel feedstock to reduce the energy emissions. Doing so requires retrofitting the kiln's firing system, as traditional fossil-fuel firing systems cannot handle the increased residue formation of biofuel combustion. Secondly, a newly developed technology allows replacing parts of the limestone with clay, which emits less CO_2 when heated. Producing calcined clay cement requires significant upgrades to the rotary kiln and control systems, as the kiln's operating temperature and rotation speed must be adjusted. Thirdly, remaining emissions can be captured and stored underground (CCUS).

Figure 3 plots the emission savings by technology type, illustrating that all levers play a substantial role in the investment strategies of firms.[4] Importantly, all levers are individually scalable, and I observe them employed to varying degrees. For example, the percentage of biofuel a firing system can accommodate varies, and the cost increases as the desired biofuel tolerance increases. Equally, the maximum replacement rate of limestone with clay is not set and can be varied based on the level of adjustment made to the production line. Lastly, the cost of a CCUS system increases with its capture rate and becomes increasingly expensive as the CO_2 concentration in the flue gas stack is reduced by other abatement levers. Further, the individual levers can be combined. For example, I observe calcined clay cement being produced with bio- or fossil-fuel heating, and a biofuel-based production line can be equipped with a carbon capture system to capture the remaining chemical emissions. Reflecting the menu of combinable, individually scalable options available to firms during my sample period, I will model the abatement investment cost as a continuous function.

The French Cement Market. The French cement market is oligopolistic, with four major players - VICAT, LafargeHolcim, EQIOM, and Ciments Calcia - accounting for over 90% of production ca-

[4]Integrated refers to investment events, where I observe kiln upgrades where at least two of the discussed technology options are deployed at the same time.

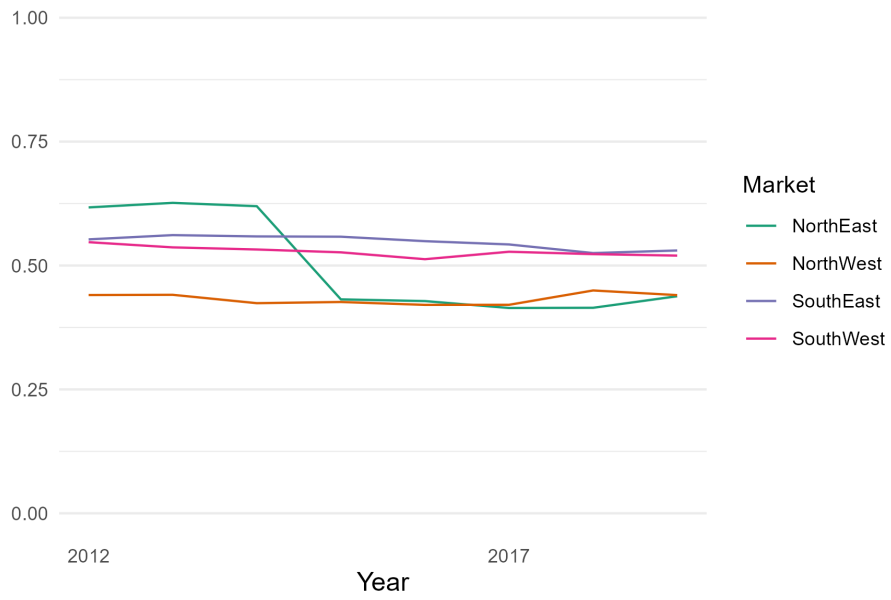


Figure 4: Herfindahl-Hirschman Index across market over time

pacity. Those four players together operate 25 integrated cement plants throughout France.[5] In my sample period, the aggregate installed capacity of integrated cement plants within France amounts to roughly 29 Mt/year.[6] Generally, competition in the cement industry is regionally segregated due to high transportation cost relative to its low value per ton (Miller & Osborne, 2014; Fowlie et al., 2016). The European Commission established in a merger review that 70% of cement is consumed within 150km of the production facility, and 90% within 250km (EC, 2014). Whilst inland transport costs prohibit extensive trade within the country, there is import competition via ship (Fowlie et al., 2016). The import share in France fluctuates slightly below 10%.

In line with the available demand data, I follow the industry organisation France Ciment and divide the French market into four regions, which are considered as independent markets in this paper. Determining which plants serve which market, I allocate plants based on their location, and assign borderline cases based on their distance to the next urban centre.[7] Figure 5 shows the location of all integrated cement plants in France, disaggregated by market allocation and ownership. For the remainder of the paper, I aggregate plants owned by the same firm within a market, and firms within each market compete as one individual agent. Their capacity constraint is derived from the sum of the capacities of the plants of a given firm in this market.

Figure 4 plots the Herfindahl-Hirschman Index for each of the markets. The HHI is typically stable over time, with the drop in the North East region explained by mandated asset divestment in the context of the Holcim-Lafarge merger.[8] The HHI values between 0.4 and 0.6 indicate significant market concentration in all markets across the sample period.

[5]Integrated cement plants include both clinker production and grinding, as opposed to facilities only grinding ready clinker into cement powder.

[6]Capacities have remained broadly fixed over time. I treat the occurring capacity changes as exogenous, but anticipated by all agents. The details are discussed in Section 5.

[7]I am currently running robustness checks in which borderline cases are assigned differently.

[8]I discuss how I deal with changes in capacity and ownership, including the merger, in the Data Section 5.

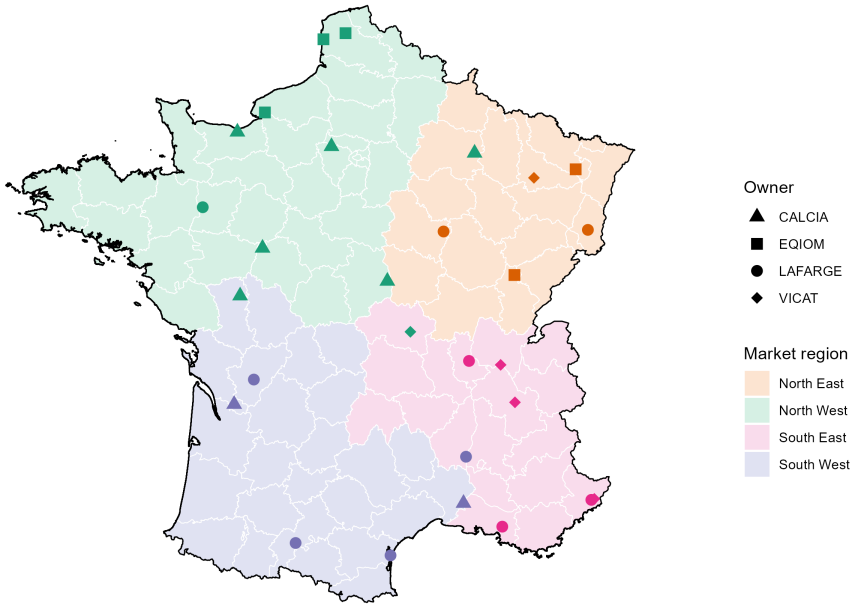


Figure 5: Regions and plant locations

The EU ETS. The European Union's cement industry has been fully regulated by the carbon price imposed by the EU ETS, an emissions-trading system, since 2005.[9] Within the ETS, regulated firms are required to submit permits for every ton of GHG-equivalent they emit. Covering both the industrial and power sectors, the EU ETS regulates the emissions of roughly 11,000 installations. Importantly, French cement installations represent a small share of covered emissions and are thus price-takers, not influencing the carbon price itself. As depicted in Figure 1, a 2016 reform of the EU ETS, combined with the post-COVID economic recovery, has driven the carbon price to increase significantly, hovering around €70-€100/ton since 2022. Figure 6 plots the share of total marginal cost represented by the EU ETS carbon price over time at an emission-intensity of 0.63.

The escalation of the carbon price since the EU ETS reform in 2016 means that, at full emission intensity, around 40% of the marginal cost is due to EU ETS compliance costs. However, as plotted in Figure 1, this sharp price escalation was accompanied by a ramp-up in low-carbon investment in the French cement industry.

Figure 7 plots the realised carbon price burdens (dashed lines) against the theoretical burden had the firms not invested (black solid line, as in Figure 7). It shows investments that have translated into substantial marginal cost savings for some firms and significant heterogeneity. The heterogeneity in investment behaviour across firms and markets resembles the key variation exploited in the dynamic estimation of this paper.

In addition to the carbon price increase, subsidy programs have been rolled out to directly support low-carbon investments. I collect detailed subsidy data from the European State Aid Transparency Database. Thus, I can explicitly account for subsidies paid out in parallel to the carbon price regulation in the dynamic estimation.

[9] Stringent regulation started to bite after the pilot phase ended in 2008.

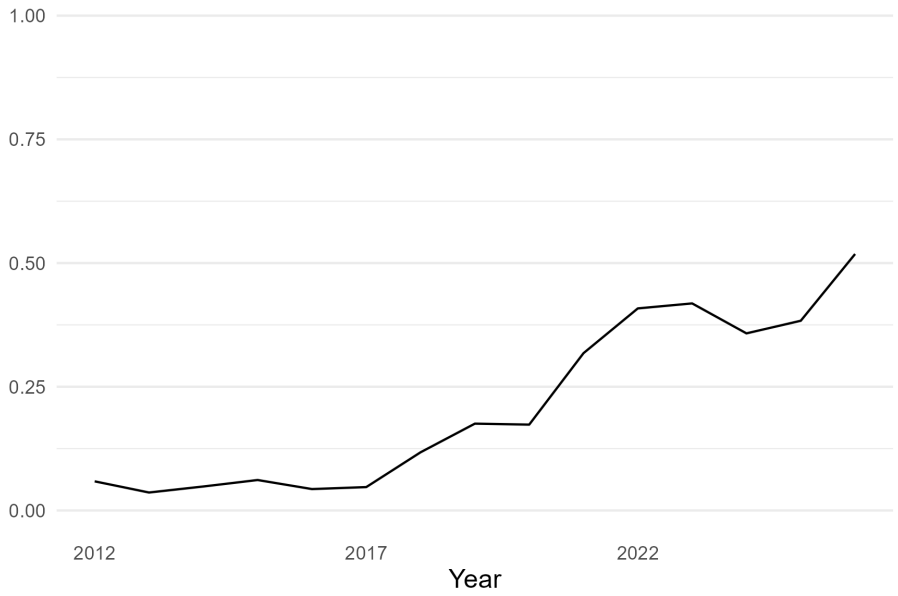


Figure 6: Share of carbon price in marginal cost

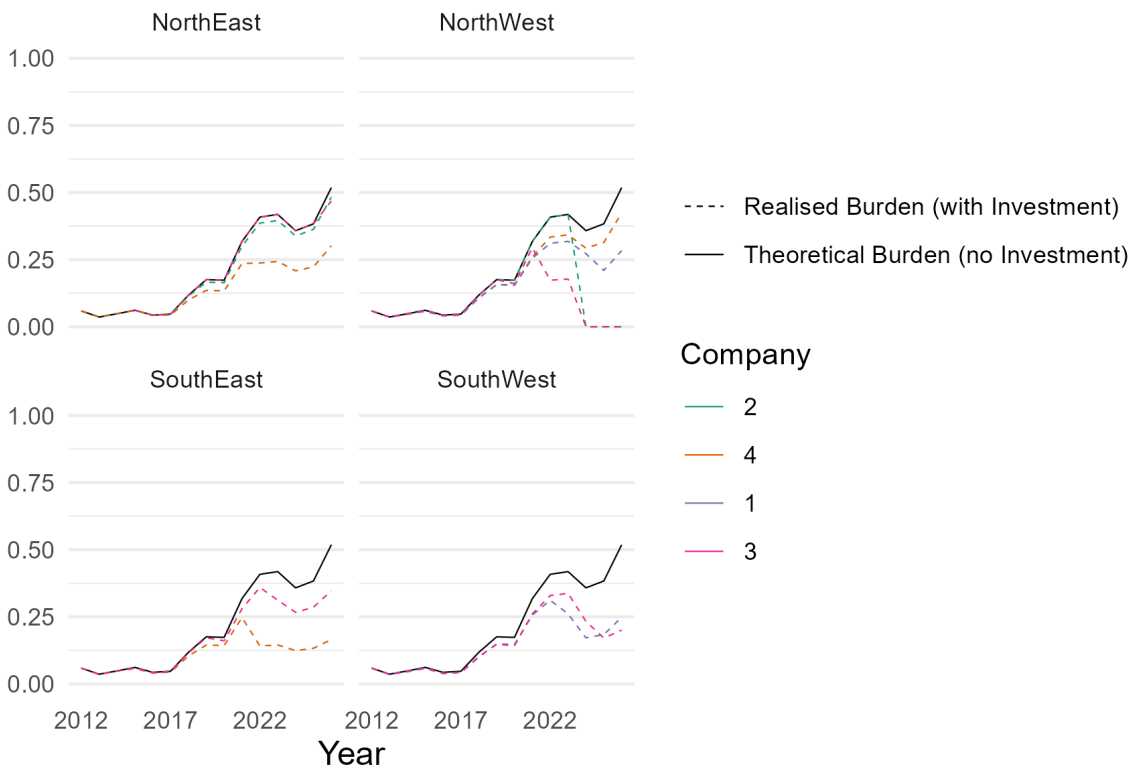


Figure 7: Share of carbon price in marginal cost with and without investment

4 Model

The model contains two parts: a static Cournot competition in every period, and a dynamic game of abatement investment. In the static Cournot model, I largely follow Ryan (2012) and Fowlie et al. (2016). Every period, firms compete in each market, facing a capacity constraint derived from aggregating the capacities of their individual plants. Firms generate revenues in the product market, incurring a cost for their emissions in addition to their production cost.

In the dynamic part of the model, firms can invest to reduce the emission intensity of their output, which alters their competitive position in the per-period game through the cost of emissions. Therefore, each market, consisting of N firms, is defined by two $N \times 1$ state vectors, s and η. s denotes the vector of production capacities, exogenously given, and η the vector of emission intensities, which evolves endogenously depending on the low-carbon investment choices made each period.

I set up an infinite-horizon model in discrete time, where β denotes the discount factor.[10] At the beginning of each period, firms receive idiosyncratic technology cost shocks and decide on their abatement investments. Then they compete in the product market, and their investment decisions mature. The state transitions are deterministic given the investment decisions.

Demand, Imports and Competitive Fringe. In each market and period, firms compete facing an isoelastic aggregate demand curve:

$$\ln Q_{m,t}(P_{m,t}; \alpha) = \alpha^0_{m,t} + \alpha^1 \ln P_{m,t} \tag{1}$$

where $Q_{m,t}$ denotes the aggregate market quantity in market m in period t; $P_{m,t}$ indicates the price, $\alpha^0_{m,t}$ the market-year intercept, and α^1 the constant demand elasticity. Further, there is an import supply curve given by:

$$\ln I_{m,t}(P_{m,t}; \alpha) = \rho^0_{m,t} + \rho^1 \ln P_{m,t} \tag{2}$$

where $I_{m,t}$ denotes the aggregate imports into market m in period t, $\rho^0_{m,t}$ the market-year intercept and ρ^1 the constant import supply elasticity. For better readability, the subscripts m, t are dropped in the following.

Whilst the four major players dominate the market, smaller actors account for approximately 10% of total capacity. To account for this, I introduce a price-taking fringe into the model.[11] Thus, the strategic players compete facing the residual demand curve, subtracting import volumes and the production of the competitive fringe from aggregate demand. Emissions implied by import volumes will be accounted for as leakage in the welfare computations.

Production and Environmental Compliance Cost. As introduced by Ryan (2012), firms face a hockey-stick production cost function, which is characterised by constant marginal cost up to a threshold of capacity utilisation, after which the marginal cost start to escalate. This cost structure reflects, for example, the extra cost incurred by foregoing maintenance, which requires an annual downtime in production, or overtime pay for night-time and weekend shifts. The production costs are:

$$C(q_i; \delta, \eta_i, \tau) = \delta_1 q_i + 1\left[q_i > \delta_2 s_i\right] \delta_3 s_i (q_i/s_i - \delta_2)^2 + \varphi(q_i, \eta_i, \tau) \tag{3}$$

[10] The results in this paper were computed with $\beta = 0.975$.

[11] A detailed list of the fringe firms, and their date of entry into the French market, is provided in Appendix A, Table 8.

where q_i is the annual production by firm i, and s_i the firm's annual capacity. δ_1 denotes the constant marginal production cost, $\delta_2 \in (0,1)$ the capacity utilisation threshold after which costs start to escalate, and δ_3 determines the slope of the cost escalation above the threshold. $\varphi(q_i, \eta_i, \tau)$ denotes the environmental compliance costs, which are part of the production cost. A firm's taxable total emissions are determined by its output q_i, emission intensity η_i, and the carbon price τ. The functional form of φ will depend on the counterfactual policy. In the case of a carbon tax, it reads as:

$$\varphi(q_i, \eta_i, \tau) = \tau q_i \eta_i \tag{4}$$

Every period, firms choose production quantities to maximise their per-period profit. The corresponding optimisation problem is given by:

$$\max_{q_i} \Pi_i(q_i; \delta, \eta, \tau) = P(Q)q_i - C(q_i; \delta, \eta_i, \tau) \tag{5}$$

In the following, $\Pi_i^*(q_i; \delta, \eta, \tau)$ will denote the optimised per-period profit, which is the function of the vector of production cost parameters δ and the vector of emission intensities of all firms, η.

Incentives for Emission Reduction. There are two channels through which a reduction in emission intensity affects a firm's per-period payoff when facing environmental regulations. Here, I will illustrate the mechanisms for the case of a carbon tax.[12]. Firstly, there is a direct cost saving for the existing production quantity. For a reduction of the emission intensity by $\Delta\eta$, the direct cost savings are:

$$-\frac{\Delta C}{\Delta \eta_i} = -\tau \Delta \eta_i q_i \tag{6}$$

Secondly, recognising that emission intensity enters the firm's marginal cost, it becomes evident that abatement investment decisions also have strategic implications. An equivalent reduction in the emission intensity reduces the marginal cost of a firm by:

$$-\frac{\Delta MC}{\Delta \eta_i} = -\tau \Delta \eta_i \tag{7}$$

Thereby, a firm alters its competitive position within the static game through its dynamic investment decision: By reducing its emission intensity, the firm gains market share relative to competitors, altering its equilibrium profits in subsequent periods.

For a marginal reduction in emission intensity, one can illustrate the gain in optimal profit for one firm using the envelope theorem (adapted from Yi (1999)):

$$-\frac{\partial \Pi_i^*}{\partial \eta_i} = q_i \left[\underbrace{\tau}_{\substack{\text{Direct} \\ \text{saving}}} - \underbrace{P'(Q)\frac{d(Q-q_i)}{d\eta_i}}_{\substack{\text{Price hike due} \\ \text{to comp. contraction}}} \right] \tag{8}$$

The strategic component of the reduction is represented in the term $\frac{d(Q-q)}{d\eta_i}$, which denotes market price increase due to the contraction of competitors' output in response to a marginal cost reduction by firm i. This is not a full characterisation of a firm's investment incentives. Firstly, the benefits of

[12]The mechanisms are equivalent for any regulation where the compliance cost increases in emissions.

expanding its own production volumes drop out through the envelope theorem. Secondly, competing firms can equally invest and may alter their own dynamic strategy in response, as a reduction in a firm's emission intensity alters the investment incentives of all competitors. The full model provides a quantitative assessment of both the gain from non-marginal investments and the investment equilibrium among competitors, using the estimated dynamic investment and competition model.

Welfare. Welfare in the model is the total sum of consumer surplus, firm profits and carbon revenue generated, minus damages from emissions and investment costs.

$$\begin{aligned} W = CS\Big(Q^{Dom} + Imports\Big) + \sum_i Profits + CarbonRevenue \\ - \sum_i InvCost - Damages\Big(E^{Dom} + Leakage\Big) \end{aligned} \tag{9}$$

I assume an emission damage function with constant marginal damages:

$$Damages(E) = \gamma E \tag{10}$$

where γ denotes the constant social cost of carbon, and E denotes total emissions:

$$E = E^{Dom} + Leakage = \sum_{j=1}^{N} \eta_j q_j + I\eta_{Imports} \tag{11}$$

The counterfactual accounts for the leakage due to imports. I assume the import fringe has a constant emissions intensity equal to the domestic baseline without investments ($\eta_{Imports} = 0.63$, see below). The leakage emissions enter into the welfare through the damages from emissions additively, and the damages are a linear function of total domestic and foreign emissions.

Private vs. Social Incentives to Reduce Emission Intensity. To understand whether private and social incentives to invest in lower-emission technologies align, consider the gain of a marginal reduction in emission intensity. Firms will set the cost to achieve this reduction against their private profit, whilst from a social perspective, the reduction is beneficial if the costs are lower than the social gain. Thus, comparing the social and private gains from marginal reductions in emission intensity shows whether firms might choose not to invest when it would be socially desirable, or invest even when the total welfare gain would not justify the cost.

Let Δ denote the change induced by an investment that lowers emission intensity. Using the welfare definition in (9), disregarding carbon revenue, the welfare effect of a marginal intensity reduction, to be compared to the investment cost, can be decomposed as

$$\Delta W = \sum_{j=1}^{N} \Delta\Pi_j + \Delta CS - \gamma\, \Delta E \tag{12}$$

where $\Delta\Pi_i$ is the change in firm i's optimised per-period profit from (5), ΔCS is the change in consumer surplus, and γ is the constant marginal damage parameter.

There are two distinct cases to consider: (i) the gain for firm i if only firm i invests, and (ii) the gain for each firm if all firms invest. When a single firm invests, its private return is captured by $\Delta\Pi_i$.

Welfare additionally accounts for rivals' profits, consumer surplus, and the externality. Rearranging (12) yields

$$\Delta W - \Delta\Pi_i = \underbrace{\sum_{j\neq i} \Delta\Pi_j}_{\substack{<0 \\ \text{(Business stealing)}}} + \underbrace{\Delta CS}_{\substack{>0 \\ \text{(Market expansion)}}} + \underbrace{\left(-\gamma\,\Delta E\right)}_{\substack{>0 \text{ since } \Delta E<0 \\ \text{(Emissions reduction)}}} \tag{13}$$

The term $\sum_{j\neq i} \Delta\Pi_j$ captures business stealing: a unilateral reduction in η_i lowers firm i's effective marginal cost, expands its equilibrium output, and induces a contraction of competitors.[13] Hence, part of the private gain $\Delta\Pi_i$ reflects a redistribution from rivals rather than net surplus creation. Thus, even though firm i does not consider the change in consumer surplus and emission damages, through the business-stealing channel, its unilateral private incentives can exceed the welfare gain. Thus, if a firm expects to be the only one to reduce its emission intensity, it is unclear whether its incentives exceed the social benefits of doing so.[14]

This is not the case for a common reduction in intensities (e.g., $\eta_j \mapsto \eta_j - \Delta$ for all j). In that case, the business-stealing channel disappears, and the welfare change is

$$\Delta W = \Delta\Pi + \Delta CS - \gamma\,\Delta E, \qquad \Delta\Pi \equiv \sum_{j=1}^{N} \Delta\Pi_j. \tag{14}$$

For a common investment, the wedge between welfare and aggregate profits is therefore

$$\Delta W - \Delta\Pi = \underbrace{\Delta CS}_{\substack{>0 \\ \text{(Market expansion)}}} + \underbrace{\left(-\gamma\,\Delta E\right)}_{\substack{>0 \\ \text{(Emissions reduction)}}}. \tag{15}$$

Since firms do not internalize ΔCS, common private incentives are generally weaker than the social incentives. Therefore, whether in equilibrium, social incentives are lower than private incentives depends crucially on the expectations and strategies of all firms. My full model will account for these mechanisms, where the gain for each firm is derived from expected profit gain, given an expectation of other firms' strategies.

Abatement Cost Function. To reduce their emission intensity, firms incur investment costs. Accounting for the fact that firms face a menu of scalable and combinable abatement levers (see Section 3), I model firms as facing a continuous abatement investment cost function:

$$InvCost_i(\theta_i|\lambda, \varepsilon_i) = s_i\theta_i\left(\lambda_1 + \lambda_2\theta_i + \lambda_3\varepsilon_i\right) \tag{16}$$

where $\theta_i \in [0, 100]$ represents the percentage reduction in emission intensity compared to a baseline intensity $\overline{\eta}$, denoting the intensity absent any abatement efforts. I set the baseline intensity to 0.63 tons of CO_2-equivalent per ton of cement (ADEME, 2023). Thus, there is a one-to-one mapping of the realised emission intensity, η_i and the investment level of the firm, θ_i:

$$\eta_i = \frac{100 - \theta_i}{100} * \overline{\eta} \tag{17}$$

[13] Under some non-standard conditions on demand and competitive position, aggregate emissions can increase when a firm reduces its intensity. For the specifications in this paper, this is not the case.

[14] I abstract from the rare, theoretical case where a reduction in emission intensity by one firm can increase aggregate emissions.

I will use these terms interchangeably. Following the newly introduced framework by Gowrisankaran and Schmidt-Dengler (2024), I introduce cost shocks $\varepsilon \sim N(0,1)$, affecting the marginal cost of investment, scaled by λ_3. The model allows for matching the observed variation in investment behaviour across similar conditions by introducing ex-ante stochasticity into the optimal choice. Thereby, I can accommodate the fact that all firms invest, and the variation in my setting lies in the extent of their investment. This is in contrast to settings where the main variation lies in which firm invests, in which case a stochastic shock to a fixed cost of investment is the appropriate modelling choice (Ryan, 2012; Fowlie et al., 2016).

During the sample period, firms invest repeatedly in each plant, adopting low-emission modernisations in multiple steps. In the model, I handle this in the following way: starting from any current state $\theta_i < 100$, firms that invest to reduce their emission intensity further to $\theta_i' > \theta_i$ incur the cost of the additional step, computed as:

$$InvCost(\theta_i'|\theta_i) = InvCost_i(\theta_i', \lambda, \varepsilon_i) - InvCost_i(\theta_i, \lambda, \varepsilon_i) \tag{18}$$

Given data constraints, I am unable to identify firm-level heterogeneity in the investment cost function. Conceptually, any heterogeneity is captured in the cost shocks, which are drawn from a joint distribution. All technology levers observed here are, in principle, mature technologies available to every firm. Thus, I hope the assumption of equal costs for these technologies across firms, net of the random shock, is not too stark.

Value Function. In every period, firms choose their emission intensity for the next period. Taking all previous elements together, they optimise the following value function:

$$\max_{\theta_i'} V_i(\theta; \delta, \lambda, \varepsilon_i, \sigma(\theta)) = \Pi_i^*(q_i; \delta, \eta, \tau) - InvCost(\theta_i'|\theta_i) + \beta \int E_{\varepsilon_i}\left[V_i(\theta'; \delta, \lambda)\right] dP(\theta'_{-i}, \sigma(\theta)) \tag{19}$$

where $P(\theta'_{-i}, \sigma(\theta))$ is the probability distribution of competitor states, a function of the starting state vector and their strategies σ.

Equilibrium Computation - Static Cournot. Given capacities, production cost, and the vector of emission intensities, I numerically solve for the Cournot equilibrium and, thereby, equilibrium profits of each firm at every point in the state space. To jointly solve the firms' first-order conditions, I use the stochastic MCMC-optimisation routine of Chernozhukov and Hong (2003), as implemented by Ryan (2012), which applies an adaptive random-walk Metropolis–Hastings search (Metropolis-within-Gibbs) to find production quantities that minimise the deviation between marginal revenue and marginal cost. The competitive fringe is explicitly modelled as a price-taking agent with a capacity constraint, producing until its marginal cost equals the market price. The price and its derivative (and thereby marginal revenue) for a given level of domestic production are computed by solving for the import volume that clears the market, given the market-level demand and import supply functions. Thus, I subtract imports that would result from a given level of the domestic output from the aggregate market demand. Thereby, I numerically implement the residual demand curve that domestic firms face.

Equilibrium Computation - Dynamic Investment. Dynamically, I compute the equilibrium investment strategies by solving for the Markov-perfect Nash equilibrium of the investment game. This

means that the strategies are Markov, thus unobservable to competitors within any given period, and depend on the current state vector. In this case, the equilibrium strategies, $\sigma_i^*(\theta)$, mapping states into choice probabilities for investment levels, $\sigma : \theta \to P(\theta')$, satisfy:

$$V_i(\theta; \varepsilon_i, \sigma_i^*(\theta), \sigma_{-i}(\theta)) \geq V_i(\theta; \varepsilon_i, \sigma_i'(\theta), \sigma_{-i}(\theta)) \;\; \forall i, \sigma_i'(\theta) \tag{20}$$

The results in this version of the paper were computed on a state space with a grid step of size 10%.[15] Thus, I explicitly compute the per-period payoff and option value for all possible combinations of firms' emission intensities in 10% increments. However, to allow for a more granular choice per firm, I approximate the value function and policies at off-grid states using multilinear (hypercubic) interpolation. For any state that lies between grid points, I identify the 2^N surrounding grid nodes and compute the interpolated value as the multilinear combination of these corner values, using weights equal to the relative position of the off-grid point along each dimension. Details of the interpolation algorithm are provided in Appendix C. I allow firms to choose their optimal emission-intensity reductions in increments of 5%. Each firm's investment choice is represented as a discretised set of available investment bins (Gowrisankaran & Schmidt-Dengler, 2024). Each bin corresponds to all investment levels between two allowed investment increments, e.g., in my case, $(0, 5]$, $(5, 10]$, etc.[16]

Firm strategies in this setting are represented as vectors of conditional choice probabilities (CCPs) over available investment bins. The CCP assigned to a bin represents the likelihood that a firm will choose an investment level within the interval defined by that bin, e.g., $(0, 5]$, $(5, 10]$. I use the algorithm provided by Gowrisankaran and Schmidt-Dengler (2024) to rapidly compute the vector of CCPs for a given vector of payoff values. The algorithm exploits increasing differences in firms' value functions to define thresholds for the cost shock between which a given investment bin is optimal. Intuitively, if an investment level of 20 is optimal for a given cost shock draw, the level of 25 will never be optimal if the cost shock realises an even higher cost. Similarly, if the level of 15 is optimal for a given draw, an even more favourable cost draw would never make 10 the optimal choice. This property is satisfied in this setting, and more generally in any model where (1) the cost of the investment is increasing with the extent of the investment, e.g., the extent of a capacity expansion, and (2) the gross payoff (not subtracting investment cost) is increasing in the investment amount. The likelihood of choosing a bin is obtained as the probability mass of the cost shock distribution falling within the interval of shock realisations for which this bin is the optimal choice. I refer the interested reader to the detailed treatment of the framework in the original paper for further elaborations.

I solve for the equilibrium by embedding this algorithm into value-function iterations, representing all combinations of emission intensities amongst the competing firms on the discretised state space. The value for a firm of being in a given state is comprised of the per-period profit it generates, defined by its relative competitive positions, an option value derived from the option to invest to reach states with lower emission intensity, and an expectation of the competitors' behaviour, shaping the expected competitive position in the following periods. The value function iteration involves initialising the value of each state as a function of the per-period profit, and computing each agent's optimal strategy given these values and an initial expectation of competitor behaviour. The values are then updated using the agents' new policy functions. The process is repeated until convergence in the values and policies is reached. Importantly, the equilibrium does not define realised investments

[15] The North West market with four players is an exception, where I set the grid to 20% to manage computational cost.

[16] Naturally, the allowed investment level when reaching 100%, which represents reaching zero emission intensity.

and industry paths, but rather a vector of choice probabilities for each situation in the game that firms can find themselves in.[17]

5 Data

In this section, I discuss the various data sources I use to estimate the model and simulate the policy scenarios. Table 4 in Appendix A provides descriptive statistics for the main variables, and more detailed information on the data sources is given in Table 5.

Demand and Imports. Demand and import data come from France Ciment, the official industry organisation of the French cement industry, for the years 2012 to 2019 (InfoCiment, 2019). Consumption data are disaggregated by region and year (the regions are depicted in Figure 5). For prices, I use administrative data on firms' self-reported prices. For each market and year, I compute the market price as the weighted average of the firm-level prices from all firms active in that market. Given the high transportation costs, imports play a minor role in the French cement market, accounting for slightly below 10% of aggregate demand. Nonetheless, I account for them to ensure accuracy. Further, in counterfactuals with high carbon prices, the import margin can affect the quantitative results. Since imports are available only at the national level, I allocate them to regions based on market size and market tightness. The resulting import penetration rates per year and market are plotted in Appendix A, Figure 21.

Capacities. Each firm runs multiple production facilities throughout France. Data on production capacities per plant come from International Cement Review by CemNet, and I cross-check them with publicly available data from the Global Energy Monitor.[18] Using the nearest urban centre, I allocate each firm's production facilities to one of the four markets. Figure 5 shows the location, ownership, and market allocation for every integrated cement plant in France. Firms within a market then compete with the aggregate capacity of all their plants serving the market. Capacity changes are not endogenous in my model, as there are insufficient observations to estimate the associated costs credibly. Still, I hand-collected a dataset of capacity changes, firm closures, and ownership changes since 2012. All instances of plant closures, transformations into grinding centres, and capacity changes are reported in Appendix A, Tables 6 and 7.[19] During my sample period, there was also a merger in 2015 of two major global players in the construction materials industry, both active in France. Holcim acquired Lafarge to become LafargeHolcim. Using the detailed decision of the European Commission on this merger, I precisely track the ownership changes and mandated divestment of assets in France (EC, 2014). In my computations, I will account for these changes, accurately representing the market situation in every year. Thereby, in the dynamic estimation, I implicitly assume firms anticipated them correctly. The counterfactuals are computed with the ownership structure and capacities as of 2025.

[17] The discretisation can lead to numerical issues. For some combinations of grid step size and model parameters, non-convergence can occur with a small number of states. In these cases, I implement a tie-breaking rule with a 0.5 update dampener for the last five iterations, until a maximum of 250 iterations is reached. In this paper, this issue rarely occurs; however, it may become more frequent when replicating the paper at different grid resolutions.

[18] In the cases where there is a discrepancy, I assume the capacity that is closest to the production volume.

[19] Grinding plants are production facilities that grind up clinker from integrated cement plants, typically within a firm. For each firm and market, I add on the grinding capacity to its production capacity when the capacity utilisation would otherwise exceed 100%.

Production. Data on production in physical quantity from 2012 to 2019 are administrative data provided by INSEE, but are only available at the firm level. To compute production volumes per firm and market, I infer plant-level production volumes from EU ETS Transaction Log (EUTL) emissions data. I assume a baseline emission factor of 0.63 and use realised abatement investment data to derive each plant's emission intensity. This allows me to approximate the production quantity from yearly realised emissions. To sanity-check the inferred production quantities, I aggregate back to the firm level and compare them with the administrative data on firm-level production. Figure 8 reports a histogram of the factor of inferred over observed production per firm and year.

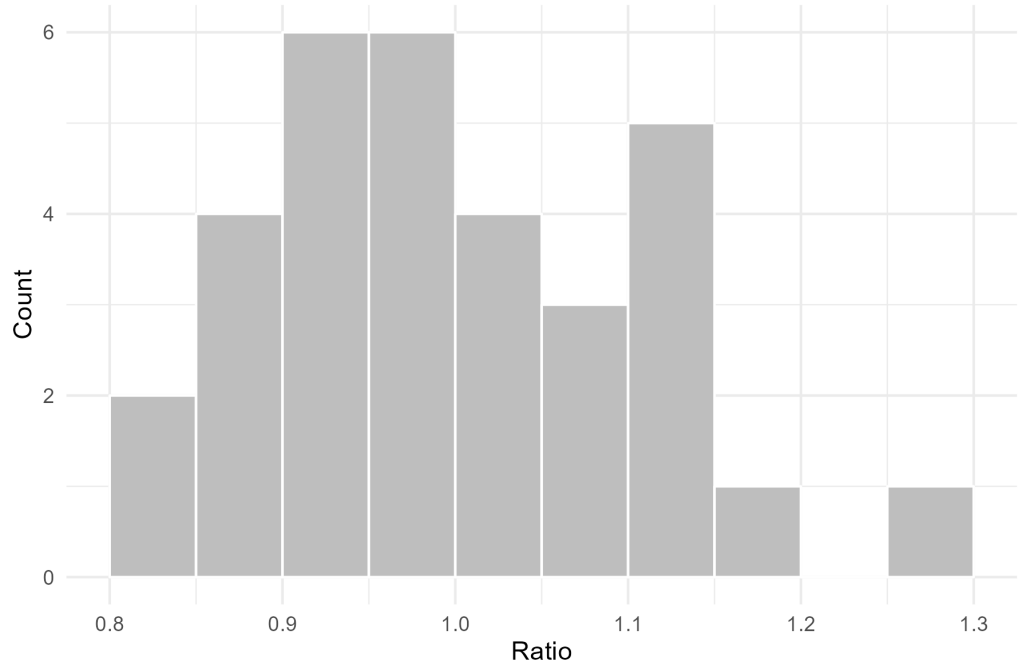


Figure 8: Factor - Inferred versus observed production

The outlier outside an error margin of 20% is the observation for the firm established through Holcim's divestments as part of the merger between Lafarge and Holcim in 2015, in the year of its creation. The residual variation most likely arises from slight differences in cement types produced, unobserved efficiency differences and in clinker imports.[20] Figure 27 in Appendix B shows the plot by firm, indicating limited, but present heterogeneity across firms. They are likely driven by variations in product composition and by reliance on clinker imports for grinding, which I cannot observe at the firm level.

Abatement Investment. I construct a novel dataset that tracks low-carbon investments in the French cement industry from 2012 to 2025. I combine administrative data with hand-collected firm announcements. Disaggregated by year and plant, the administrative data, taken from the ANTIPOL dataset by INSEE, report the monetary values associated with engineering costs of investments dedicated to reducing GHG emissions from 2012 to 2022. For the model's economic cost estimation, when costs include opportunity costs, overhead, and financing costs, monetary spending needs to be converted into reductions in emission intensity.

[20] Portland cement, highly standardised, accounts for over 85% of production volumes. Nonetheless, there is a share of other types of grey cement, whose emission intensity might not fully match.

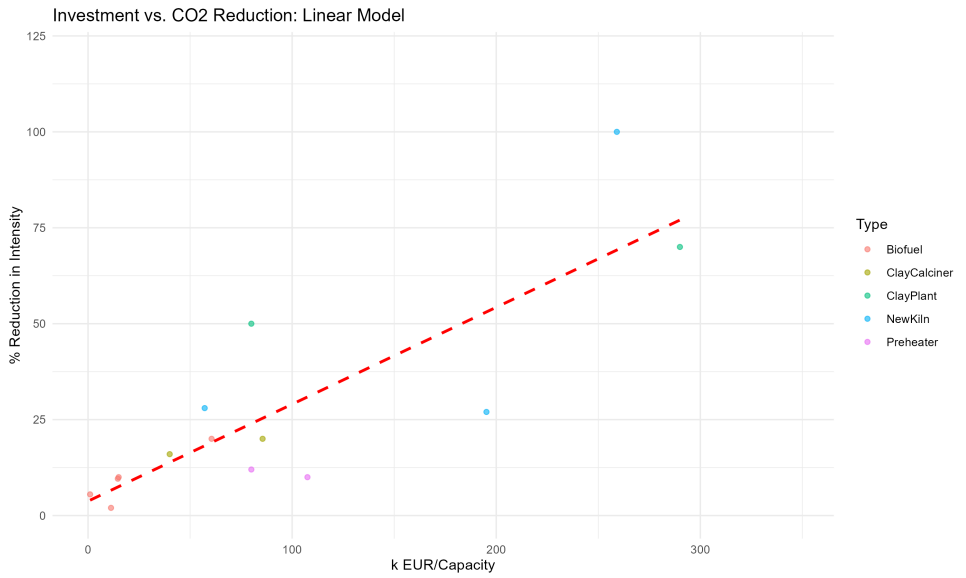


Figure 9: Engineering Cost - Abatement Investment Cost Curve

To do so, I hand-collect a dataset of company announcements that report both cost and emission reductions from low-carbon investments. I use these data points to construct an engineering cost curve, which enables me to translate reported spending into reductions in emission intensity. Figure 9 depicts the relationship I exploit between euros spent per capacity and percentage reductions in emission intensity. Since the administrative data are currently only available up to 2022, I further complement this data with the latest investments announced to capture meaningful variation in carbon prices and investments for the recent years from 2023 to 2025 (see Figure 1). The complete list of investments I use to complement the ANTIPOL data is reported in Appendix A, Table 9. Furthermore, CCUS projects are typically reported in terms of their capture capacity. To compute the savings rate in emission intensity, I derive the reduction in emission intensity implied by the capture capacity and the plants' 2019 emissions, accounting for announced capacity changes.[21]

Subsidies. As part of national and EU-wide programs, firms within my sample period have received subsidies for low-carbon investments. In the EU, subsidies paid to firms have to be reported in the European State Aid Transparency Database. I collect the data on the investment subsidies received by French cement firms, and as for the abatement investments, hand-collect the most recently announced subsidies.

6 Estimation: Method and Results

This section discusses the estimation strategies separately for the model's static and dynamic components. I proceed by discussing the static estimations first, since the results feed into the dynamic estimations by determining the static payoffs in the dynamic optimisation. All objective functions in the estimations are minimised using a Covariance Matrix Adaptation Evolution Strategy.[22]

[21] I chose 2019 to use pre-COVID figures. As data become available, I will update this using the newest data.

[22] For an explanation of the method, I refer the reader to the review by Hansen (2016). I use the Apache Commons Math Java implementation, version 3.6.1 (Apache Software Foundation, 2016).

6.1 Estimation - Static Parameters

Demand and Imports. Demand and import elasticities are taken from the literature at this stage (Fowlie et al., 2016). Concretely, I set:

$$\alpha^1 = -2 \tag{21}$$

$$\rho^1 = 0.5 \tag{22}$$

Using the data on demand and prices, in combination with these elasticities, I fit year-market-specific intercepts to construct the demand and import supply curves for every year and market.

Production Cost. To estimate the production cost parameters, I follow the approach in Ryan (2012), and minimise two equally weighted moments, evaluated per firm, year and market: (1) the squared deviations from observed production quantities and the model predictions and (2) the squared difference between marginal revenue and marginal cost at the observed quantities.[23] I minimise the moments by estimating three parameters, the set of which is denoted by the 3×1 vector δ: the constant marginal cost, δ_1, the capacity utilisation threshold δ_2, and the escalation of cost above the threshold, δ_3.

In the estimation, I need to account for the existing regulations firms faced during the sample period. Firstly, firms faced the EU ETS carbon price. Specifically, I compute the carbon price component of the firm's cost based on its emission intensity and the average EUA price in this year. Thus, I can isolate the constant part of the marginal cost curve, net of the compliance costs incurred under the EU ETS during the sample period. Secondly, a further complication is the presence of free-allocations rules during the sample period. The production estimation sample (2012-2019) largely falls into Phase III (2013-2020) of the EU ETS. In previous periods, free allocation was entirely based on historical output (HAL) and thus did not distort contemporaneous production decisions. Phase III introduced Activity Level Correction Factors (ACLFs), which established thresholds for the current-to-historical output ratio, above which the free allocation was updated. Notably, as long as firms produce at least 50% of their average historic production volume (see Branger et al. (2015) for details). This can, in principle, create incentives to alter production decisions in order to increase the allocated free allowances. At the installation level, I consistently observe production volumes that are significantly higher than 50% of historical levels. Thus, within my sample, they are comfortably above the 50% utilisation threshold. Therefore, the incentive distortion, induced by the ACLFs around the threshold, is limited in this context, and the choice of production level is unlikely to be strongly affected by the free allowances.

The results are reported in Table 1.[24]

MC (δ_1)	71.74 EUR/ton
Utilisation threshold (δ_2)	76.66%
Utilisation cost (δ_3)	0.505 EUR

Table 1: Estimated cost function parameters

The utilisation cost implies that, after exceeding 77% of capacity utilisation, the marginal cost of firms increases by €1.1 ($2\delta_3$) with every additional percentage of utilisation. The estimated con-

[23] The estimator has equally been employed by Fowlie et al. (2016).

[24] Standard errors for the estimates are work in progress.

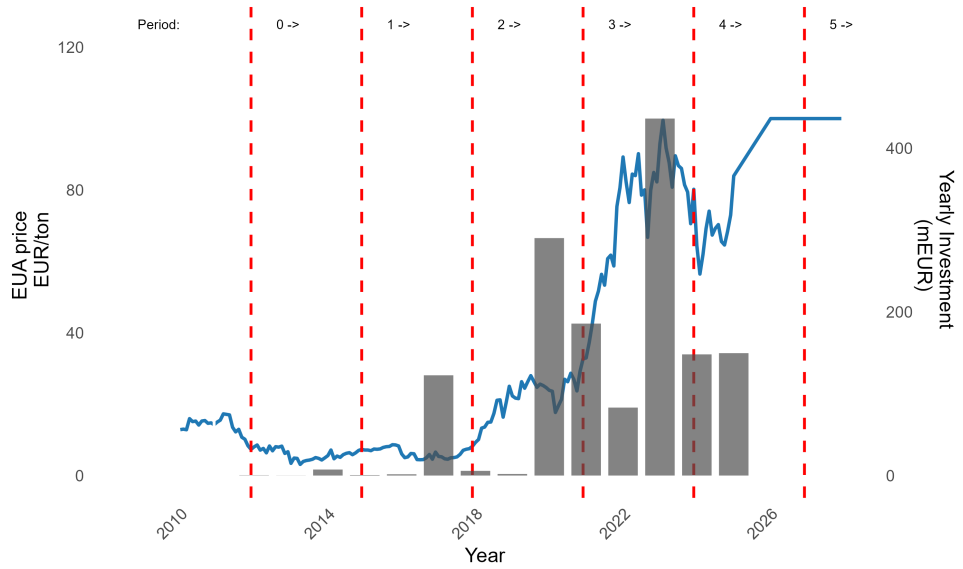


Figure 10: Carbon pricing regimes and investment periods

stant marginal cost of €72/ton, and the average price for a ton of cement during my sample period is €109/ton. Combined with the fact that firms incur costs for exceeding the utilisation threshold and for their emissions, this implies an operating margin hovering between 30-40%. I plot the profit margins implied by my estimates per firm and market over time in Appendix B, Figures 23 and 22. Given the high capital intensity, this mark-up is within the expected range and close to the estimates by Fowlie et al. (2016). I report the model's moment fit for the production quantity per market and per firm and the market-level price in Appendix B, Figures 24, 25, and 26.

6.2 Estimation - Dynamic Parameters

To estimate the parameters governing the dynamic abatement investment decision, I employ a full-solution nested fixed-point GMM approach. This involves two steps: first, to avoid the computational burden of solving the full dynamic game for every guess in the parameter search, I estimate reduced-form competitor responses. Individual firms optimise against this response curve, turning the dynamic game into a single-agent optimisation instead of a full dynamic game. Then, I determine optimal investment decisions for firms throughout the sample period by solving the model for the infinite-horizon continuation value in a final period and backward-inducting the optimal strategy in each period to the beginning of the study period (Gowrisankaran et al., 2024). This allows me to generate the model-implied investment behaviour for each sample period, which I can match against observed investment moments.

Identifying Variation and Investment Periods. The low-carbon investments studied here involve significant construction times. Therefore, estimating the model using yearly decisions rationalises investments by focusing on short-term variations in carbon prices and competitors, rather than on more meaningful longer-term variations. Given different carbon pricing regimes, I estimate the model using three-year investment periods, with carbon prices over these periods shown in Figure 10.

I assume that firms face a long-term price expectation of €125/ton, which is realised from 2026 onwards. Then, I solve for the infinite continuation value at €125/ton and backward-induct the op-

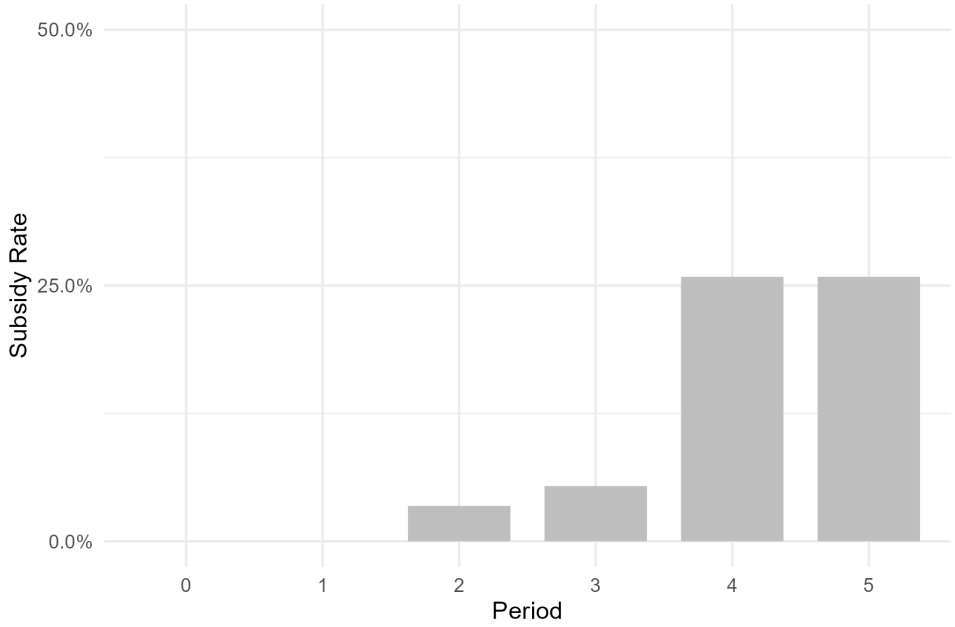


Figure 11: Subsidy rate per Period

timal investment strategies to 2012 in three-year increments. The identifying variation comes from firms' investment responses to variations in the carbon price. Concretely, the estimates will rationalise which firms invested how much in response to the EU ETS allowance price increases since 2012. I control for the subsidy programs at the national and EU levels, rolled out in parallel with the escalation of the EUA price. Using data from the European State Aid Transparency Database, I compute the subsidy rate received by firms in each period, defined as the sum of subsidies received divided by the sum of investments across firms, and assume that all firms in that period were eligible for that rate. I assume the latest observed subsidy rate, period 4, is the one firms face from 2026 onward (period 5). Figure 11 plots the observed subsidy rates for each investment period.[25]

Reduced-form Competitor Responses. Solving the full dynamic game every time for an extensive search among candidate parameters is computationally expensive. To address this issue, I estimate the reduced-form competitor responses, defined as the percentage reduction in emission intensity (0 to 100) realised by competitors, as a function of the period, market share, and the emission intensity at which the firm starts the period. In the estimation, firms optimise, expecting this response curve from their competitors, turning the dynamic game into a single-agent problem. These are much quicker to compute and allow a more extensive search of candidate parameter values for estimation. Notice there is no observed competitor strategy for the future infinite horizon continuation value in period 5. I let firms optimise against the competitor response, as in the last observed period, period 4. Specifically, I estimate the following linear specification:

$$\theta_{i,(t+1)} = \beta_1 \,\text{MarketShare}_{i,t} + \sum_{k=0}^{4} \gamma_k \, 1\{\text{Period}_t = k\} + \beta_2 \,\text{stateExAnte}_{i,t} + \epsilon_{i,t}. \tag{23}$$

As a consequence of the low observation count, the identified reduced-form competitor responses, reported in Table 2, are largely statistically insignificant. Still, the responses show the expected signs.

[25] I only include subsidies received for specific abatement projects, not R& D funding.

Losing market share reduces the incentives to invest: first, the direct cost savings are lower, and second, isoelastic demand curves imply declining industry marginal revenue. Therefore, catching up to competitors in a context where they have already invested, resulting in higher aggregate production, is not as beneficial as expanding production when your competitors still produce less due to high intensity. Further, an already lower emission intensity reduces a firm's propensity to invest even further, reflecting the escalating abatement costs. Period fixed effects capture the varying investment levels across periods. Reflecting the low number of observations, I do not observe firms in all possi-

	Dependent variable:
	Competitor State $\in (0, 100)$
Market Share $\in (0, 1)$	10.111
	(11.335)
Period 0	-2.471
	(6.683)
Period 1	-0.789
	(6.648)
Period 2	3.021
	(6.611)
Period 3	14.337
	(6.770)**
Period 4	29.528
	(8.510)***
exAnte $\in (0, 100)$	0.880
	(0.233)***
Observations	55
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 2: Competitor policy functions

ble states of the game and cannot fully identify the conditional choice probabilities empirically. Thus, I employ a full-solution estimator, where the exact representation of competitor responses in the estimation is a second-order concern. The firms' behaviour I set against the data is derived from a full solution of the value function over time. This is in contrast to the estimator proposed by Bajari et al. (2007), where the identification relies more strongly on a precise estimation of the reduced-form policy functions.

Abatement Investment Cost Function. Using these competitor response functions, I fully solve the resulting single-agent optimisation problem for each firm and parameter candidate, and compare the model outcomes to data moments. More specifically, I estimate the parameters of the abatement cost function by minimising two equally weighted moments, evaluated per market and period: (1) the squared deviation of the implied mean of investment and observed investment, and (2) the squared deviation of the implied against the observed standard deviation of investment between firms within a market. The first moment, matching the expected average level of investment, primarily identifies λ_1 and λ_2. The dispersion in investment choices not explained through structural model fundamentals

is captured in the second moment, which allows me to identify λ_3. The results are reported in Table 3. As a reference point, these estimates imply an expected cost of €7.5bn to reduce the emission inten-

	λ_1	λ_2	λ_3
Estimate	3.73	0.084	4.429

Table 3: Estimates of abatement cost function parameters

sity of the entire French cement industry by 50%. For comparison, techno-economic assessments by the French environmental agency ADEME estimate the cost for the same reduction at €3.3 bn (Mari et al., 2021). My estimates, unlike the techno-economic assessments, include opportunity costs associated with project overhead and downtime. As a back-of-the-envelope calculation, my estimates correspond to an implicit opportunity cost roughly equivalent to four months of downtime, given an annual industry turnover of $\approx$ €12.5 billion (in 2024).

7 Counterfactuals

To evaluate counterfactual policy scenarios, I solve for the dynamic game's equilibrium and simulate the expected industry path forward, assuming every firm starts at full emission intensity. At this stage of the project, the counterfactual policies include different levels of carbon taxes and technology cost subsidies.[26]

Equilibrium Computation. Here, I briefly recap the method used to compute the equilibrium of the dynamic investment game. For more details on the equilibrium computation, I refer the reader back to Section 4. To solve for the Markov-perfect Nash equilibrium of the game, I perform value function iterations on a discretised grid, approximating the value function and best responses using hypercubic interpolation. I use the algorithm provided by Gowrisankaran and Schmidt-Dengler (2024) to integrate over the cost shocks and compute best responses quickly. The algorithm exploits the increasing-difference property of the cost function to determine unique thresholds for cost shocks, such that each grid level is optimal within its corresponding threshold. This allows to solve for the equilibrium as the fixed point of value functions and best responses. The results are equilibrium choice probabilities at each grid point, which allow computation of the equilibrium value of being in each state, integrating out the equilibrium expectations of competitor actions and future cost shocks.

Forward Simulation. Having computed the equilibrium values along the grid, I can forward simulate the industry path by performing Monte Carlo simulations on the technology cost shocks. Letting the industry start at full emission intensity of 0.63 (see ADEME (2023), discussed in Section 3), I draw a technology cost shock for each firm and period. Given the draw and the equilibrium values at each state, I can determine each firm's optimal decision and update its state accordingly. I forward simulate 30 investment periods and plot the net present value of the outcomes (Fowlie et al., 2016).

[26] Other scenarios, such as a carbon border adjustment or efficiency benchmark pricing, are work in progress.

Results. In this section, I show the outcomes determining the welfare described in (9) for different carbon tax and technology subsidy levels. I contrast (1) of the full dynamic model (Full Model) with (2) the static oligopoly outcomes (Static - Oligopoly).[27] Thereby, I compare my model's policy recommendations with static analyses, such as Buchanan (1969), and illustrate the relevance of the dynamic investment margin for policy recommendations. All production capacities and market constellations for the simulations reflect the situation of the French cement industry in 2025. The plotted values are aggregated across markets. Emission damages are set at €125/ton.

Consumer surplus, determined by the aggregate market quantity, is plotted in Figure 12. A carbon price suppresses output by increasing the marginal cost of production. As the carbon price increases, the production levels drop in all cases. However, as firms in the dynamic model can ease the carbon tax burden by investing, the quantity is suppressed less sharply. A given tax level represents a lower addition to the marginal cost, as green technology adoption has decreased the taxable emissions per unit of output. Therefore, the economic burden due to output suppression is reduced by the firm's ability to react along a margin, in addition to output adjustments.

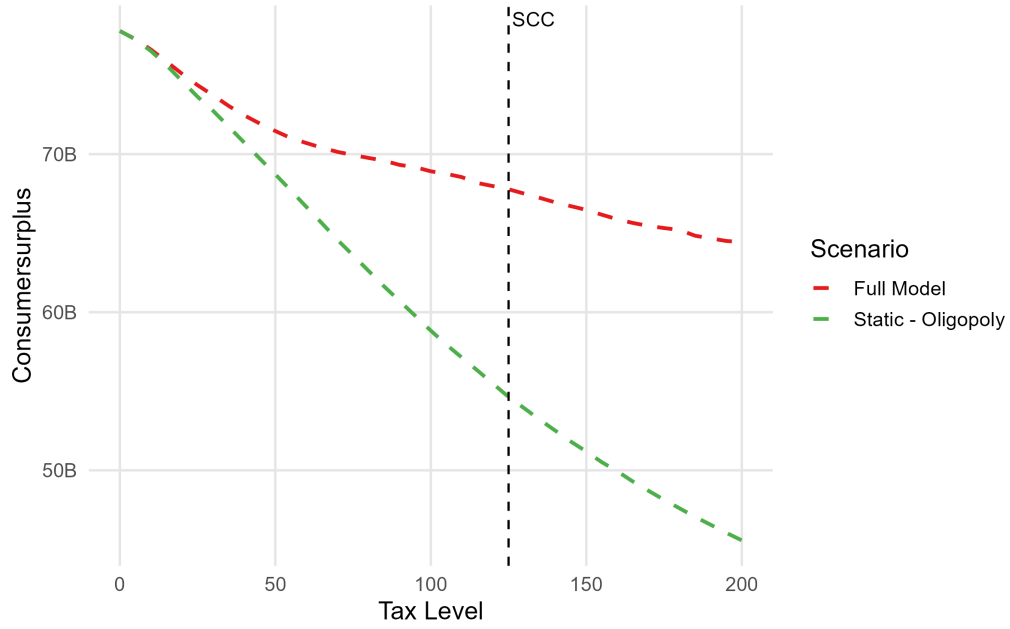


Figure 12: Net present value: Consumer surplus

Intuitively, emission damages (Figure 13) are lowered when firms can invest to reduce their emission intensity. As the carbon tax increases, they decrease more strongly in the dynamic model. Importantly, the emission damages are significantly lower even though, as described above, the production quantity in the complete dynamic model remains significantly higher than in the static setting. This illustrates well how green technology adoption benefits welfare in the economy: by decreasing the emission intensity of production, more quantity can be produced, boosting consumer surplus, without incurring excessive emission damages.

[27] The model structure implies the possibility of negative investment cost for extreme draws and low investment levels. To ensure consistency between the dynamic and static settings, I derive the development in the static oligopoly setting using the exact forward simulation used for the dynamic model, setting the investment gain to zero. This is purely for accuracy, with no implications for the results or policy.

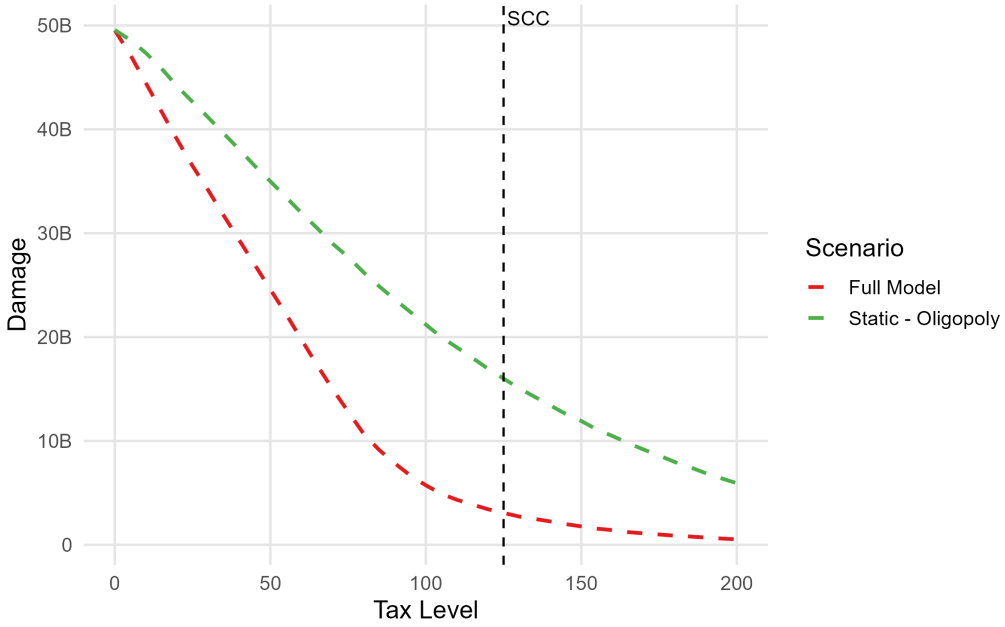


Figure 13: Net present value: Damages

However, firms incur higher investment costs to achieve this. Figure 14 shows that the firms' investment choice increases sharply and levels off as it becomes increasingly hard to incentivise even further investments at the pointy end of the abatement investment curve. Interestingly, an ever-increasing carbon price can even choke the investments, which slightly decline for high carbon tax levels above €150/ton. The reason is that very high carbon prices turn the investment increasingly into an all-or-nothing game. When the carbon price is high enough, there is no point in investing to remain slightly more emission-intensive than your competitor, as this still represents a substantial disadvantage. Thus, in equilibrium, firms either invest in complete decarbonisation or not at all. Thereby, investments in the middle of the distribution are prevented, and the overall equilibrium level turns out to be lower than with a lower carbon price.

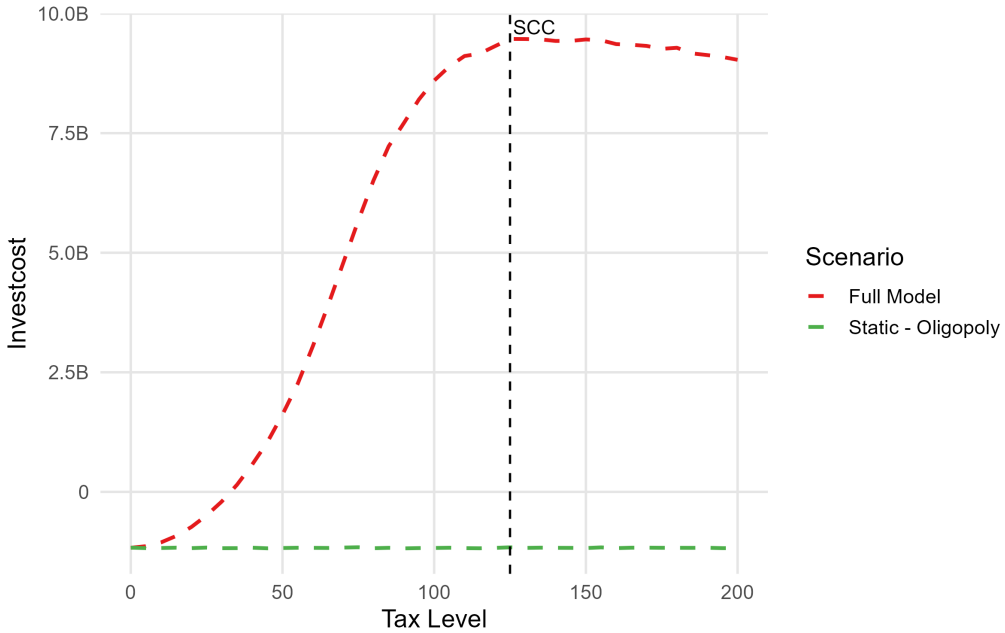


Figure 14: Net present value: Investment Cost

The option to reduce their emission intensity allows firms in the dynamic setting to keep their variable profits high, whereas increasing the carbon price causes declining profits if firms do not have this option (Figure 15).

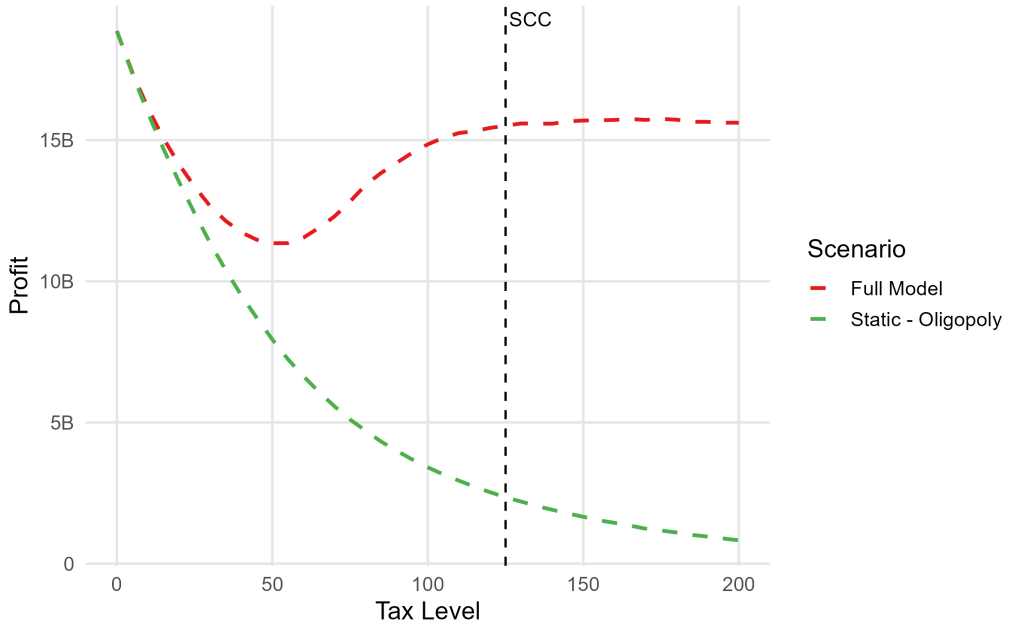


Figure 15: Net present value: Profits

The revenues generated from carbon pricing (Figure 16) follow a Laffer curve pattern. Thus, setting a very high carbon price erodes the tax base by reducing emissions sufficiently that increasing the carbon price further reduces carbon revenue. As expected, the peak occurs earlier, implying lower carbon revenues when firms can invest.

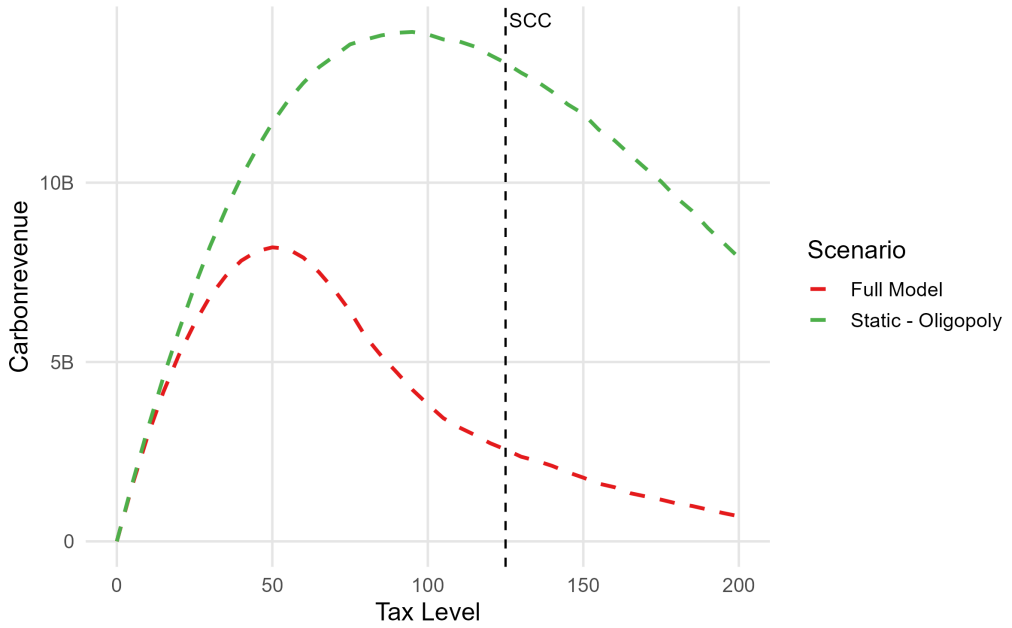


Figure 16: Net present value: Carbon Revenue

Putting all these components together yields the welfare implications of the carbon price levels, plotted in Figure 17. The static setting illustrates the Buchanan argument. The welfare-optimal carbon

price is lower than the social cost of carbon, because output is already suppressed below the socially optimal level, thereby reducing damages and consumer surplus. Reducing output further through a carbon tax is thus less beneficial. This is true in both the static and dynamic context. However, in a dynamic context, firms must also be sufficiently incentivised to invest in emission-abatement technology. Therefore, the optimal carbon tax is higher than a static analysis would suggest, but still lower than in a perfectly competitive setting.

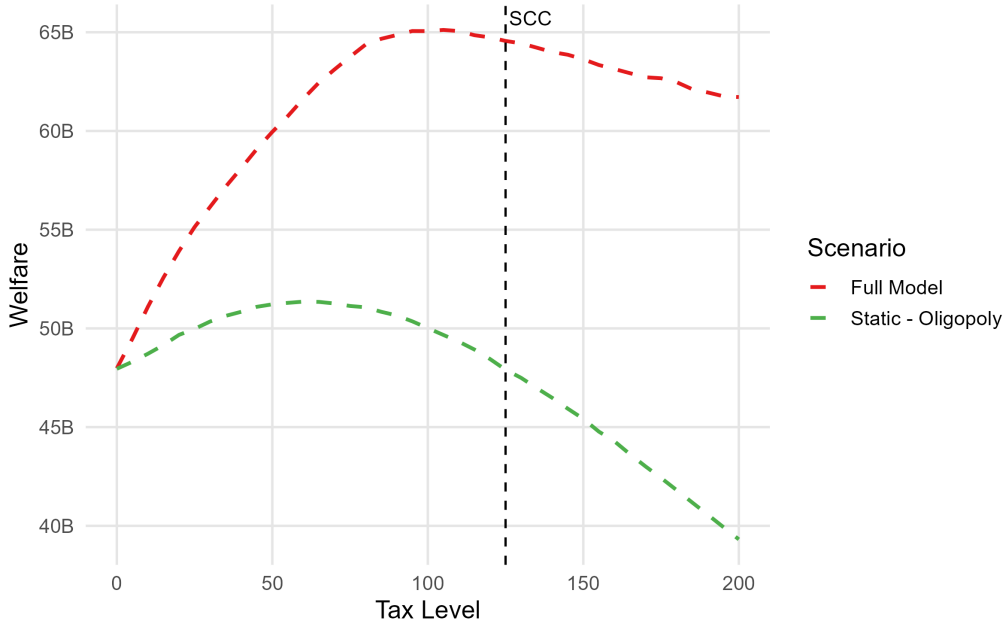


Figure 17: Net present value: Total Welfare

The results illustrate the countervailing effects when introducing carbon taxation in a concentrated industry that can invest to reduce its carbon emissions. Market power already restricts the output, lowering the welfare benefits of carbon taxation. However, when the carbon price is set too low, socially beneficial investment options are not realised, and too little is produced at too high emissions. In this setting, with a constant damage of emission of €125/ton, I estimate an optimal carbon tax of €105/ton, compared to €60/ton implied by a static oligopoly analysis. Setting the tax at the statically optimal level would result in a welfare loss of 5%.

The incentives for firms to invest in low-carbon technologies are suboptimal when they hold market power (Hoarau & Ponssard, 2025). I can confirm this in my setting by comparing the realised investment levels against the socially optimal investment level. Defining the socially optimal investment level is tricky, since the production volumes, and thus the societal gain from emission intensity reductions, depend on the tax level and product market conduct. Therefore, I compute the socially optimal investment level by computing both the value function iteration and the Monte Carlo simulation exactly as in the baseline model, but replacing the firm's profit in the value function, defined in Equation 19, by the total welfare, defined in Equation 9, generated in each point in the state-space. Doing so, I generate a benchmark investment behaviour that would result if firms invested to maximise welfare instead of profits. Setting this against the baseline model, I can isolate the firms' reduced investment incentives. The results are plotted in Figure 18. Whilst an increasing carbon tax increases investment, the plot shows that firms' investment levels remain suboptimal.

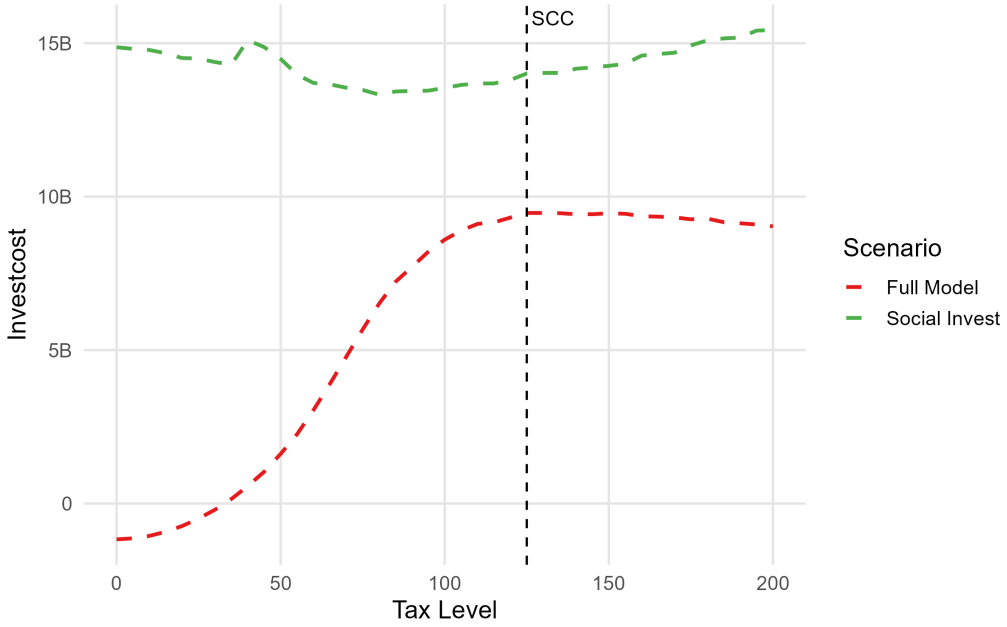


Figure 18: Net present value: Investment Cost

Therefore, a tax alone cannot achieve a socially optimal outcome. A natural consequence is the introduction of a second policy instrument. In addition to a carbon tax, I compute counterfactual outcomes, setting the carbon tax at the statically optimal level and additionally subsidising a share of firms' investment cost. Figure 19 shows the total welfare for varying technology subsidy shares in the full model. Supplementing the statically optimal carbon tax with a subsidy yields higher welfare than the optimal tax alone can achieve. The optimal technology subsidy of 55% of investment cost yields a welfare gain of $\approx 1.9\%$ above what can be achieved by setting only an optimal tax. The reason is that three externalities are present in this market: market power in the product market, emission damages, and reduced investment incentives. Interestingly, because they go in opposite directions, a carbon tax can simultaneously correct for market power and the emission externality in the static product market. However, the additional distortion due to lowered investment incentives requires further intervention such as a technology cost subsidy.

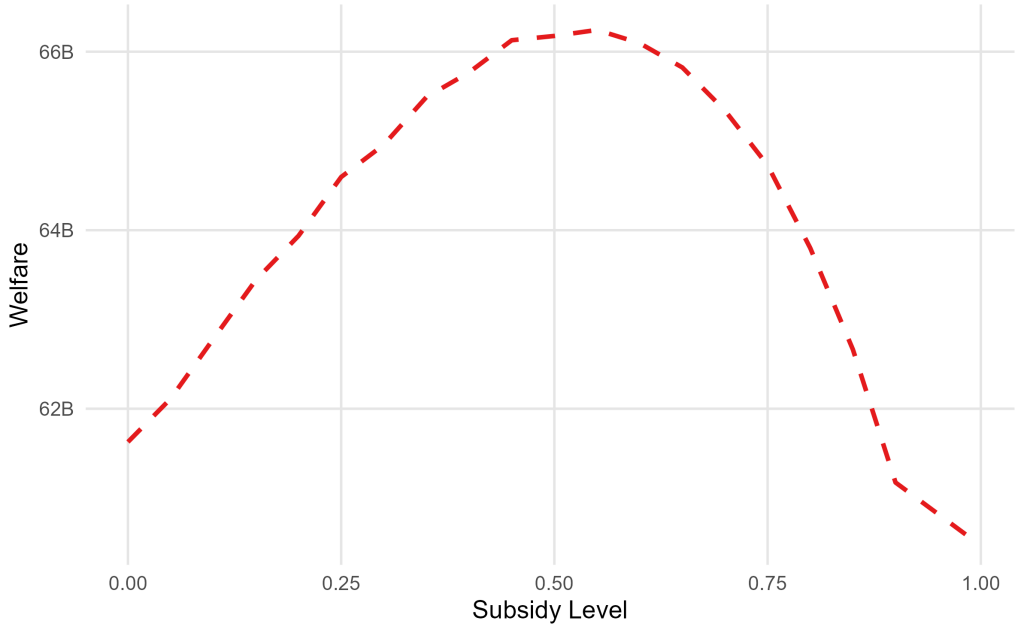


Figure 19: Net present value: Total Welfare - Technology Subsidy

A carbon tax, combined with a technology cost subsidy, can be an attractive policy mix, since revenues from the tax can be used to fund the subsidy. Figure 20 plots the net public funds for increasing technology cost subsidies. Whilst the policy is not fully self-funded, the public money required to sustain this policy is limited by the carbon revenue generated. Notice the welfare calculation in Figure 19 fully incorporate for these subsidies, using the investment cost regardless of who pays them. This implicitly assumes an equivalence between private and public cost of funds. Given the maximal welfare values obtained under both scenarios, I estimate the welfare advantage of the subsidy scenario persists up until a cost of public funds of $\approx 25\%$. Once the cost of public funds exceeds that of private funds by more than $\approx 25\%$, the welfare gain from the subsidy scenario is eliminated.[28]

[28]That is assuming only negative public funds incur cost, and positive public fund generation is weighted equally to private profits.

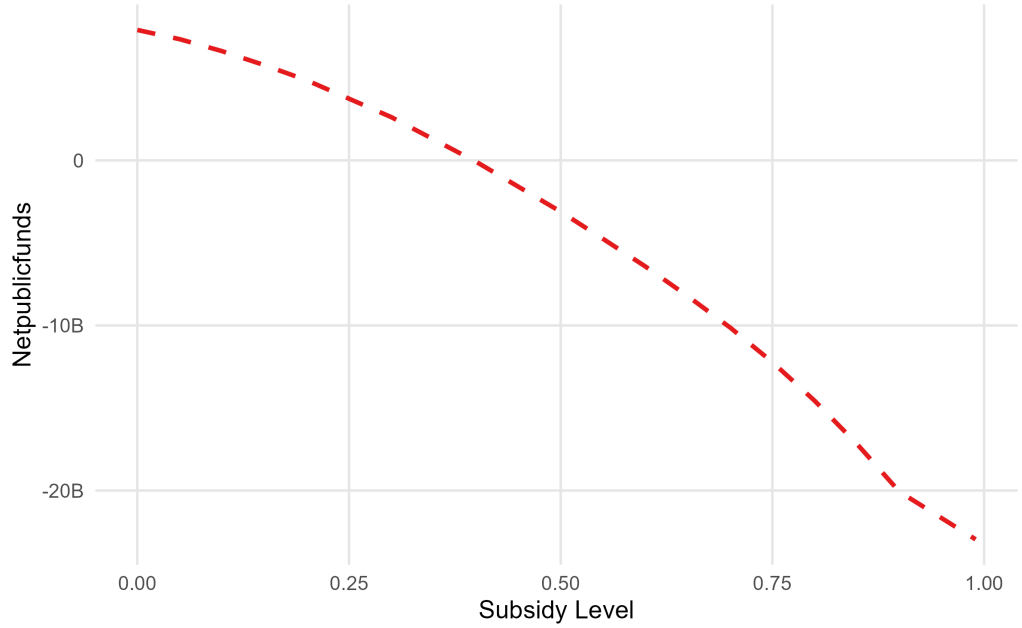


Figure 20: Net present value: Carbon Revenues - Subsidy Cost

8 Conclusion

Many countries, including the member states of the European Union, have set ambitious targets for reducing emissions over the coming decades. In the EU, as the power sector decarbonises, the industrial sector decarbonisation will pose a significant obstacle to achieving emission-reduction targets. At the same time, backlash from consumers and industry over energy prices and the regulatory burdens of climate policy is growing. Achieving emission reductions without eroding public support requires a deep understanding of the effects of climate policy measures in heavy industry sectors. Thus, to meet its net-zero targets without triggering substantial deindustrialisation, the EU climate policy is designed to reduce emissions primarily through the adoption of green production technologies rather than reductions in industrial activity. Recognising that market concentration is a key feature of the most polluting heavy industries, this paper complements the discussion of the effects of climate policy on firms by analysing the welfare dynamics arising from regulating emissions in a concentrated industry that can invest to reduce the emission intensity of its output.

This paper employs a structural dynamic oligopoly model to compute green technology adoption and the welfare implications of environmental regulation. Estimating the model with data on the French cement sector, I quantify the equilibrium investment in emission abatement using novel algorithms for estimating continuous investment choices (Gowrisankaran & Schmidt-Dengler, 2024). To my knowledge, this study is the first to fully solve and quantify the abatement investment equilibrium and the resulting industry outcomes in an oligopolistic market.

Doing so, I show that climate policy in concentrated industries needs to balance output suppression (Buchanan critique) with sufficient incentives to adopt low-emission technologies. As a consequence, the welfare-maximising carbon tax level in a concentrated industry lies between the Pigouvian level and the welfare-maximising tax implied by a static oligopoly analysis. The reason is that the economic burden of carbon prices on consumers and firms is substantially reduced if there is the option to invest in reducing the emission intensity of production. The Buchanan critique argues

that suppressed output in concentrated markets already partially internalises the emission externality, and further output contraction due to carbon prices leads to excessive loss of consumer surplus. I show that the option to adopt low-emission technology eases this constraint, as increasing the carbon price does not have an equally detrimental effect on domestic production. Still, as long as the emission intensity of production is positive, the Buchanan argument does apply, and the optimal carbon price is thus lower than the Pigouvian level. A technology cost subsidy can deliver welfare gains, but the gains are limited.

This study contributes to the policy debate on the future of climate change mitigation. If an industry has access to technologies to reduce its emission intensity, concerns about deindustrialisation and output loss are less pronounced. On the contrary, climate policy needs to incentivise firms to adopt socially beneficial low-emission technologies. Still, excessive carbon prices do hamper domestic production and even investments, and thereby consumer and producer surplus, even when low-carbon technology is available. Thus, a carbon policy that strikes a balance between the two margins can provide substantial welfare gains without imposing a significant burden on domestic firms. This is because adopting green technology enables significant emission reductions without significant output suppression. Therefore, the firms' incentives to adopt green technology in equilibrium are a major driver of the welfare implications of carbon regulations and should be thoroughly considered when introducing them. As technologies for low-carbon production become increasingly available across major emitting sectors, efficiently incentivising their adoption is critical. This paper shows that combining policy instruments, such as a carbon tax with technology subsidies, can yield a climate policy that squares the circle between reducing emissions, limiting price increases, and preventing domestic deindustrialisation.

References

Buchanan, J. (1969). External Diseconomies, Corrective Taxes, and Market Structure. *American Economic Review*, *59*(1), 174–77.

Delbono, F., & Denicolo, V. (1991). Incentives to Innovate in a Cournot Oligopoly. *The Quarterly Journal of Economics*, *106*(3), 951–961. https://doi.org/10.2307/2937936

Laffont, J.-J., & Tirole, J. (1996). Pollution permits and environmental innovation. *Journal of Public Economics*, *62*(1-2), 127–140. https://doi.org/10.1016/0047-2727(96)01576-9

Yi, S.-S. (1999). Market structure and incentives to innovate: The case of Cournot oligopoly. *Economics Letters*, *65*(3), 379–388. https://doi.org/10.1016/S0165-1765(99)00169-X

Montero, J.-P. (2002). Permits, Standards, and Technology Innovation. *Journal of Environmental Economics and Management*, *44*(1), 23–44. https://doi.org/10.1006/jeem.2001.1194

Chernozhukov, V., & Hong, H. (2003). An MCMC approach to classical estimation. *Journal of Econometrics*, *115*(2), 293–346. https://doi.org/10.1016/S0304-4076(03)00100-3

Requate, T., & Unold, W. (2003). Environmental policy incentives to adopt advanced abatement technology: *European Economic Review*, *47*(1), 125–146. https://doi.org/10.1016/S0014-2921(02)00188-5

Requate, T. (2005). Dynamic incentives by environmental policy instruments—a survey. *Ecological Economics*, *54*(2-3), 175–195. https://doi.org/10.1016/j.ecolecon.2004.12.028

Bajari, P., Benkard, C. L., & Levin, J. (2007). Estimating Dynamic Models of Imperfect Competition. *Econometrica*, *75*(5), 1331–1370. https://doi.org/10.1111/j.1468-0262.2007.00796.x

Fowlie, M. (2010). Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement. *American Economic Review*, *100*(3), 837–869. https://doi.org/10.1257/aer.100.3.837

Ryan, S. (2012). The Costs of Environmental Regulation in a Concentrated Industry. *Econometrica*, *80*(3), 1019–1061. https://doi.org/10.3982/ECTA6750

EC. (2014). Holcim / Lafarge — Commission decision. *European Commission, Directorate-General for Competition*, (Case M.7252).

Fabra, N., & Reguant, M. (2014). Pass-Through of Emissions Costs in Electricity Markets. *American Economic Review*, *104*(9), 2872–2899. https://doi.org/10.1257/aer.104.9.2872

Martin, R., De Preux, L. B., & Wagner, U. J. (2014). The impact of a carbon tax on manufacturing: Evidence from microdata. *Journal of Public Economics*, *117*, 1–14. https://doi.org/10.1016/j.jpubeco.2014.04.016

Miller, N. H., & Osborne, M. (2014). Spatial differentiation and price discrimination in the cement industry: Evidence from a structural model. *The RAND Journal of Economics*, *45*(2), 221–247. https://doi.org/10.1111/1756-2171.12049

Branger, F., Ponssard, J.-P., Sartor, O., & Sato, M. (2015). EU ETS, Free Allocations, and Activity Level Thresholds: The Devil Lies in the Details. *Journal of the Association of Environmental and Resource Economists*, *2*(3), 401–437. https://doi.org/10.1086/682343

Apache Software Foundation. (2016). *Apache Commons Math: CMAESOptimizer* (*Version 3.6.1*). Java implementation of the Covariance Matrix Adaptation Evolution Strategy (CMA-ES).

Calel, R., & Dechezleprêtre, A. (2016). Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market. *Review of Economics and Statistics*, *98*(1), 173–191. https://doi.org/10.1162/REST_a_00470

Fowlie, M., Reguant, M., & Ryan, S. P. (2016). Market-Based Emissions Regulation and Industry Dynamics. *Journal of Political Economy*, *124*(1), 249–302. https://doi.org/10.1086/684484

Hansen, N. (2016). The CMA Evolution Strategy: A Tutorial [Version Number: 2]. https://doi.org/10.48550/ARXIV.1604.00772
Other ArXiv e-prints, arXiv:1604.00772, 2016, pp.1-39.

Phaneuf, D. J., & Requate, T. (2016, December). *A Course in Environmental Economics: Theory, Policy, and Practice* (1st ed.). Cambridge University Press. https://doi.org/10.1017/9780511843839

Wagner, U., Muuls, M., & Martin, R. (2016). 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*.

Marie, M. J., Gaied, K. S., & Ajel, A. R. (2018). Wireless Network Control System for Electro Stati Precipitator in Cement Plants. *Al-Khwarizmi Engineering Journal*, *13*(4), 50–57. https://doi.org/10.22153/kej.2017.04.003

InfoCiment. (2019). L'Essentiel-Statistiques : Statistiques et chiffres clés de l'industrie cimentière française (2011–2019) [Published: Publications annuelles 2011–2019, disponibles sur Infociments].
Série de rapports statistiques publiés par Infociments couvrant la période 2011–2019.

Calel, R. (2020). Adopt or Innovate: Understanding Technological Responses to Cap-and-Trade. *American Economic Journal: Economic Policy*, *12*(3), 170–201. https://doi.org/10.1257/pol.20180135

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), 303–342. https://doi.org/10.1257/app.20180474

Koltay, G., & Lőrincz, S. (2021). Industry concentration and competition policy. *European Commission, Directorate-General for Competition*, (Competition Policy Brief 02/2021).

Mari, E., Sourisseau, S., & ADEME. (2021, November). *Sectoral Transition Plan for the French Cement Industry: Summary Report* (tech. rep.). ADEME (Agence de la transition écologique).
40 pages; online 21 December 2021.

ADEME. (2023). Feuille de route de décarbonation de la filière Ciment. *Ministére de la Transition écologique et de la Cohésion des Territoires*, (Mai 2023).

Dechezleprêtre, A., Nachtigall, D., & Venmans, F. (2023). The joint impact of the European Union emissions trading system on carbon emissions and economic performance. *Journal of Environmental Economics and Management*, *118*, 102758. https://doi.org/10.1016/j.jeem.2022.102758

EP. (2024). Greenhouse gas emissions by country and sector. *European Parliament*, (Article - 20180301STO98928).

Gowrisankaran, G., Langer, A., & Reguant, M. (2024, January). *Energy Transitions in Regulated Markets* (Working Paper No. 32088) (Series: Working Paper Series). National Bureau of Economic Research. https://doi.org/10.3386/w32088

Gowrisankaran, G., & Schmidt-Dengler, P. (2024, May). *A Computable Dynamic Oligopoly Model of Capacity Investment* (tech. rep. No. w32399). National Bureau of Economic Research. Cambridge, MA. https://doi.org/10.3386/w32399

Toyama, Y. (2024). Dynamic Incentives and Permit Market Equilibrium in Cap-and-Trade Regulation. *American Economic Journal: Microeconomics*, *16*(3), 374–423. https://doi.org/10.1257/mic.20190377

CEMBUREAU. (2025, June). *Key Facts & Figures — Production-Consumption-Emissions: Cementitious Products in the \em CEMBUREAU Region, June 2025* (tech. rep.). CEMBUREAU – The European Cement Association.

Colmer, J., Martin, R., Muûls, M., & Wagner, U. J. (2025). Does Pricing Carbon Mitigate Climate Change? Firm-Level Evidence from the European Union Emissions Trading System. *Review of Economic Studies*, *92*(3), 1625–1660. https://doi.org/10.1093/restud/rdae055

Gowrisankaran, G., Langer, A., & Zhang, W. (2025). Policy Uncertainty in the Market for Coal Electricity: The Case of Air Toxics Standards. *Journal of Political Economy*, *133*(6), 1757–1795. https://doi.org/10.1086/734779

Hoarau, Q., & Ponssard, J.-P. (2025). Imperfect Competition and the Adoption of Clean Technology: The Case of CCS in Cement. https://doi.org/10.2139/ssrn.5506619

Appendix

A Data Sources and Descriptives

Variable	Market	mean	sd	p25	p50	p75	p100
Production	NorthEast	2223.05	76.19	2183.04	2235.83	2249.05	2351.67
Production	NorthWest	6383.93	325.15	6095.96	6402.22	6574.97	6837.17
Production	SouthEast	3722.80	210.57	3604.51	3753.52	3882.25	3957.35
Production	SouthWest	4072.37	283.31	3907.46	3981.19	4197.01	4502.82
Capacity	NorthEast	2450.00	0.00	2450.00	2450.00	2450.00	2450.00
Capacity	NorthWest	7230.00	0.00	7230.00	7230.00	7230.00	7230.00
Capacity	SouthEast	5200.00	0.00	5200.00	5200.00	5200.00	5200.00
Capacity	SouthWest	3963.64	504.52	3600.00	3600.00	4600.00	4600.00
Emissions	NorthEast	1206.78	404.49	1288.54	1332.80	1368.07	1397.37
Emissions	NorthWest	3428.65	1157.97	3552.07	3622.44	3919.32	4100.73
Emissions	SouthEast	2082.83	701.58	2188.61	2276.52	2369.20	2465.54
Emissions	SouthWest	2214.55	751.51	2277.01	2407.33	2482.53	2704.36
Utilisation	NorthEast	0.91	0.03	0.89	0.91	0.92	0.96
Utilisation	NorthWest	0.88	0.04	0.84	0.89	0.91	0.95
Utilisation	SouthEast	0.72	0.04	0.69	0.72	0.75	0.76
Utilisation	SouthWest	1.00	0.09	0.96	1.01	1.08	1.10
Price	NorthEast	105.40	1.60	104.00	105.00	106.00	108.00
Price	NorthWest	112.40	4.30	109.00	111.00	115.00	120.00
Price	SouthEast	114.90	5.40	111.00	114.00	118.00	123.00
Price	SouthWest	114.50	8.50	109.00	115.00	119.00	124.00

Table 4: Summary Statistics by Variable and Market

Data Point	Source	Dataset	Level	Access
Investment	INSEE	ANTIPOL	Plant	Proprietary - via CASD
Investment	Firms/news	Hand-collected	Plant	Full dataset provided
Capacity	Cemnet	-	Plant	Proprietary
Capacity	Global Energy Monitor	Cement Tracker	Plant	Public
Emissions	EUETS.info	EUTL	Plant	Public
Production	INSEE	EAP	Firm	Proprietary - via CASD
Revenue	INSEE	EAP	Firm	Proprietary - via CASD
Demand	Infociment	LÈssentiel	Market	Public
Imports	Infociment	LÈssentiel	Market	Public

Table 5: Overview of Variables and Data Sources

Firm	Plant	Market	Event	Date
Heidelberg	Gargenville	NorthWest	Transformation	2022
Heidelberg	Beffes	NorthWest	Closure	2024
Heidelberg	VilliersauBoin	NorthWest	Closure	2024
HolcimLafarge	Frangey	NorthWest	Closure	2013
HolcimLafarge	LaCouronne	NorthWest	Transformation	2016
HolcimLafarge	LeHavre	NorthWest	Closure	2015
HolcimLafarge	Dannes	NorthWest	Closure	2014
HolcimLafarge	Contes	SouthEast	Closure	2021

Table 6: Plant Closures and Transformations

Firm	Plant	Market	Date	Capacity Expansion
ECOCEM	Dunkirk	NorthWest	2025	300
EQIOM	Lumbres	NorthWest	2023	400
Heidelberg	Airvault	NorthWest	2020	360
HolcimLafarge	StPierre	NorthWest	2022	1400
HolcimLafarge	Matres	SouthWest	2017	500

Table 7: Capacity Expansion Events

Firm	Plant	Market	Type	Capacity	Date
Lafarge	SaintNazaire	NorthWest	Grinding	600	2014
CIMAT	VracsEstuaire	NorthWest	Grinding	600	2016
Cem In EU	Tonneins	SouthWest	Grinding	240	2018
ECOCEM	Dunkirk	NorthWest	Grinding	750	2018
Cem In EU	LesValences	SouthEast	Grinding	240	2021

Table 8: Entry of fringe firms

Firm	Plant	Market	Type	M_EUR	CO2ta	CO2rate	Announced	Expansion
ECOCEM	Dunkirk	NorthWest	ClayPlant	87		70	2025	Yes
EQIOM	Lumbres	NorthWest	NewKiln	285		100	2023	Yes
Heidelberg	Airvault	NorthWest	NewKiln	285		27	2020	Yes
Heidelberg	Airvault	NorthWest	CCUS		1000		2024	No
Heidelberg	Beaucaire	SouthWest	HeatRecovery	7			2022	No
Heidelberg	Beaucaire	SouthWest	Preheater	86		10	2023	No
Heidelberg	Bussac-Foret	SouthWest	Biofuel	46		20	2022	No
Heidelberg	Bussac-Foret	SouthWest	ClayCalciner	65		20	2023	No
Heidelberg	Couvrot	NorthWest	Preheater	80		12	2024	No
HolcimLafarge	LaMalle	SouthEast	Clay	6		31	2022	No
HolcimLafarge	LaMalle	SouthEast	Biofuel	17,5	48	9,6	2022	No
HolcimLafarge	LeTeil	SouthEast	CCUS		200		2023	No
HolcimLafarge	Matres	SouthWest	Biofuel	1	8	2	2016	No
HolcimLafarge	Matres	SouthWest	NewKiln	120		28	2017	Yes
HolcimLafarge	Matres	SouthWest	CCUS		700		2024	No
HolcimLafarge	Matres	SouthWest	Biofuel	11			2024	No
HolcimLafarge	Nouvelle	SouthWest	Biofuel	6			2021	No
HolcimLafarge	StPierre	NorthWest	ClayPlant	40		50	2022	Yes
HolcimLafarge	StPierre	NorthWest	CCUS		1000		2023	No
HolcimLafarge	StPierre	NorthWest	Biofuel	13			2024	No
VICAT	Crechy	SouthWest	Biofuel	4,5		2	2018	No
VICAT	Montalieu	SouthEast	Biofuel	30	40	10	2021	No
VICAT	Montalieu	SouthEast	CCUS		500		2021	No
VICAT	Xeuilly	NorthEast	Biofuel	1	9,85	3	2020	No
VICAT	Xeuilly	NorthEast	ClayCalciner	40	48,5	16	2021	No

Table 9: Investment Events Summary (Sorted by Firm and Plant)

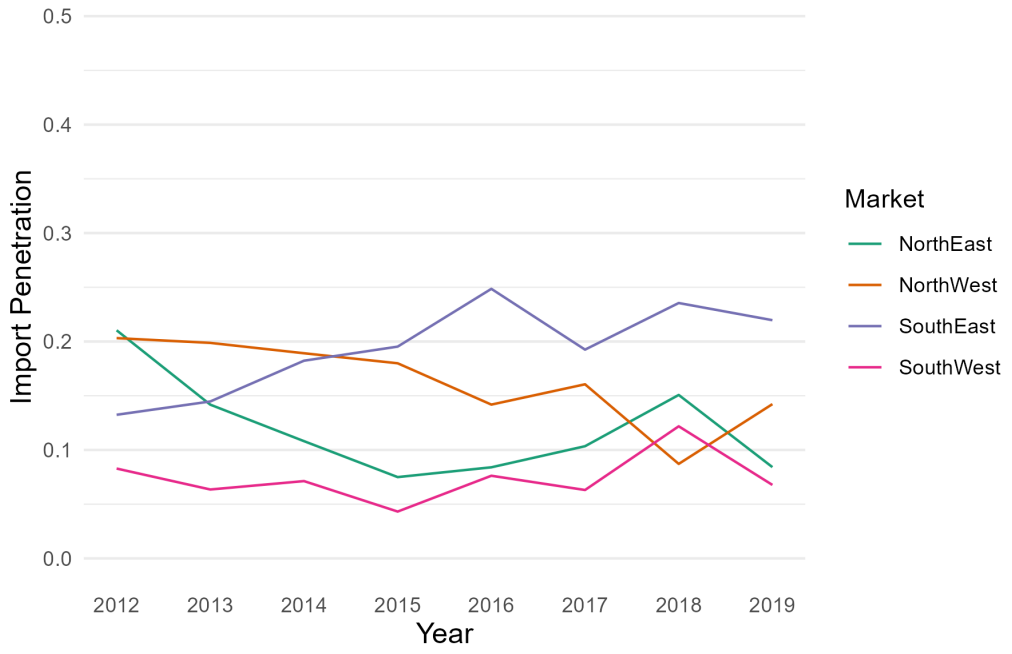


Figure 21: Import share by year and market

B Estimation

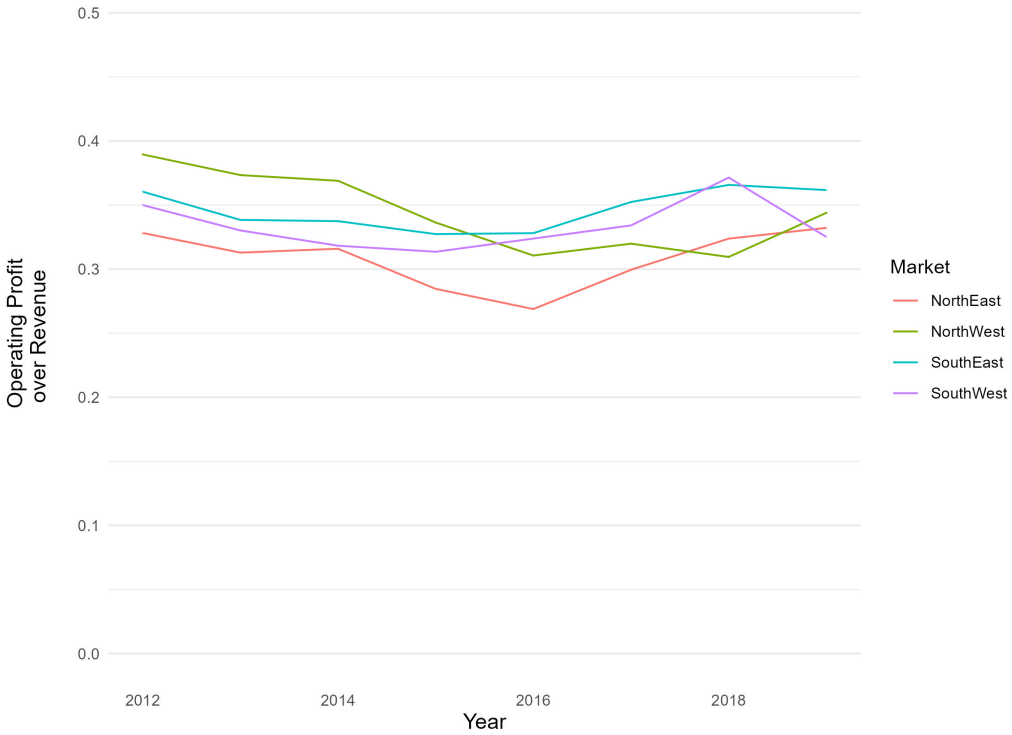


Figure 22: Moment Prediction - Operating Margin per Market

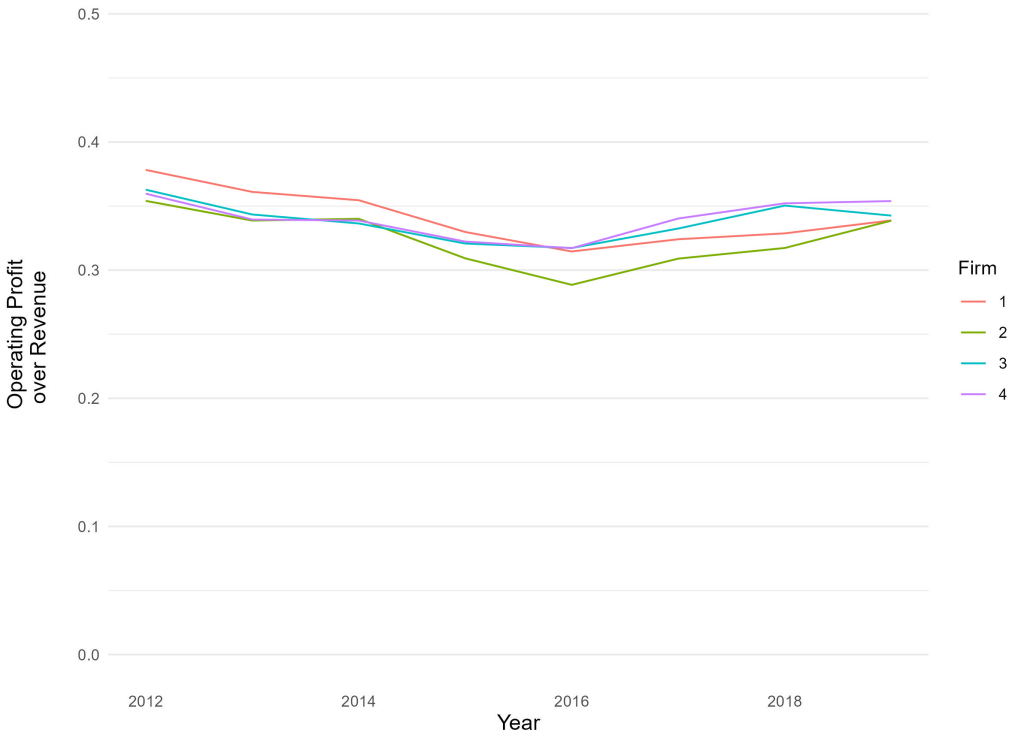


Figure 23: Moment Prediction - Operating Margin per Firm

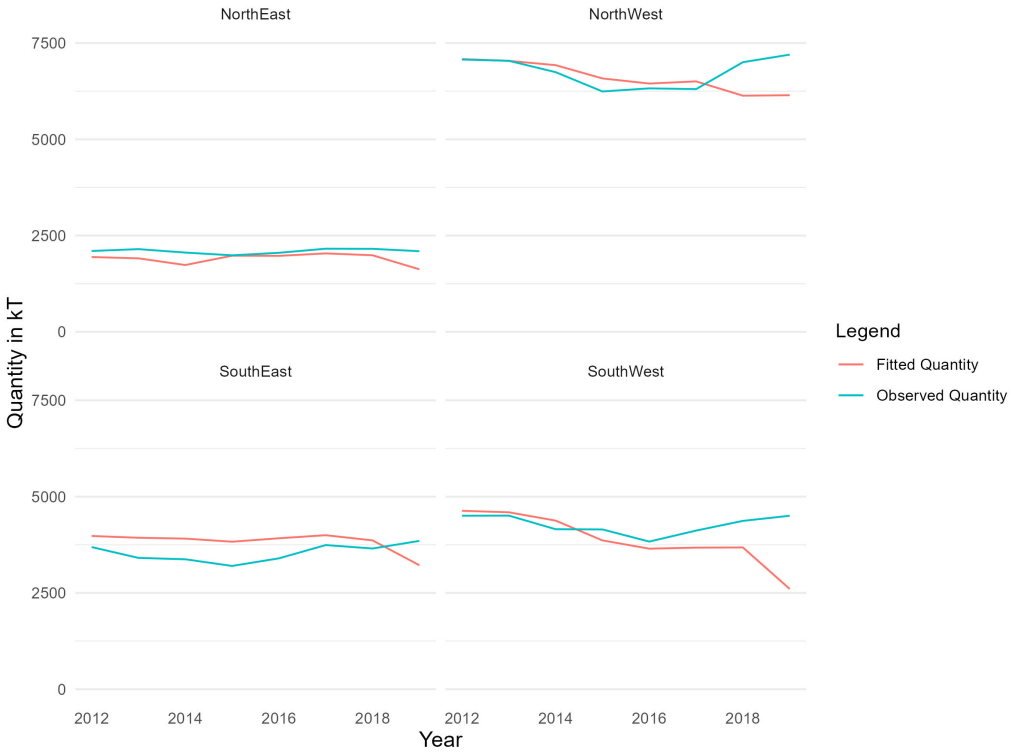


Figure 24: Moment Fit - Market level production

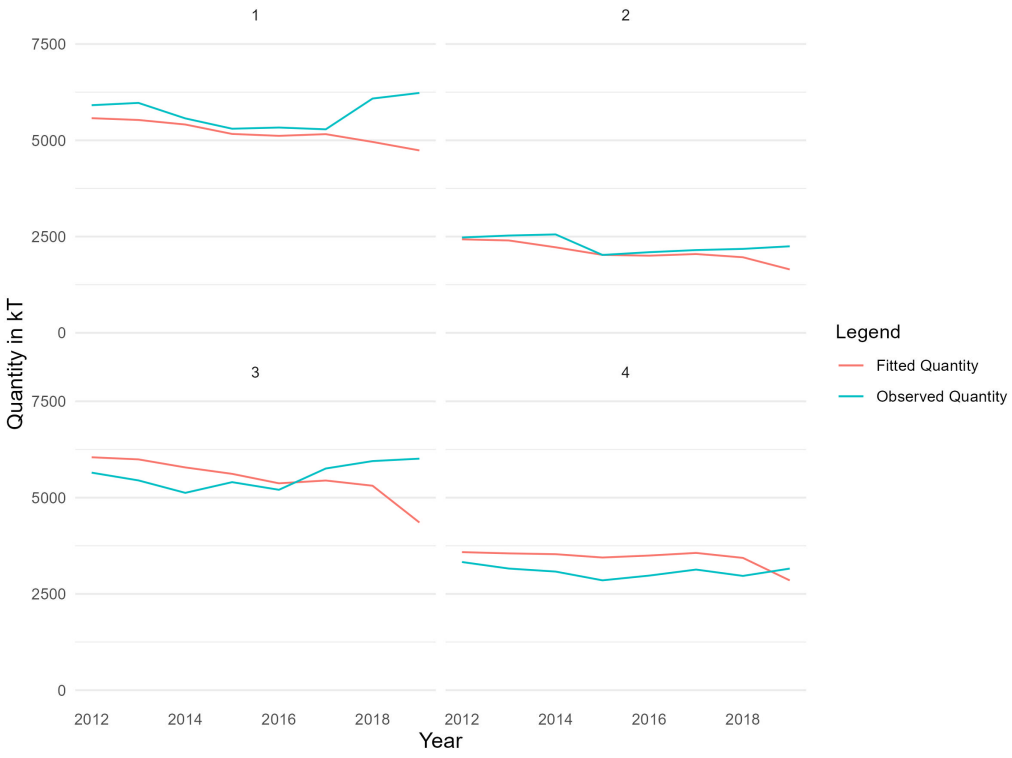


Figure 25: Moment Fit - Firm level production

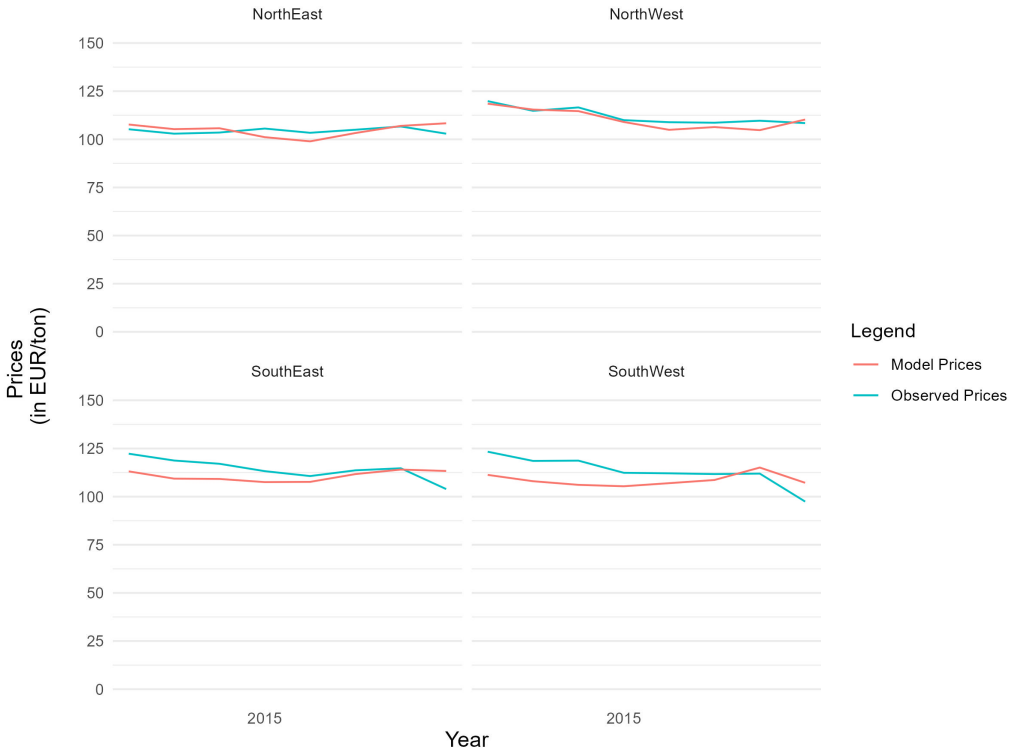


Figure 26: Moment Fit - Firm level production

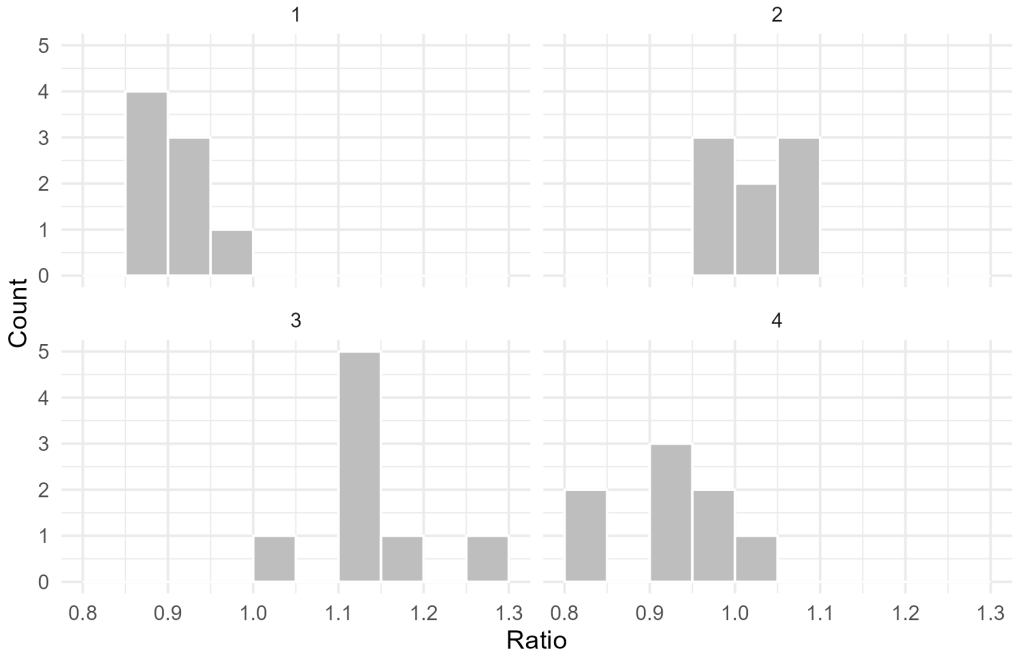


Figure 27: Factor - Inferred versus observed production

C Value Function and Equilibrium Computation

Interpolation of the Value Function. To approximate continuous investment decisions, I need to evaluate the continuation value of firms in off-grid states. To do so, I approximate the value function using hypercubic interpolation. I illustrate the details of this method here. A concrete 2x2-example is provided at the end of this section.

The algorithm for the interpolation works as follows: Let the state vector be $x^N = (x_1, \ldots, x_n) \in \mathbb{R}^N$, where N denotes the number of dimensions of the state space (in our case, N is the number of firms per market, and x^N denotes the vector of emission intensities). For an off-grid state vector x^N, I define lower and upper grid points surrounding x^N for each firm as x_n^- and x_n^+. Then, one can determine the position of any point within the state space relative to each firm's surrounding corner states. For firm n, this reads as:

$$t_n = \frac{x_n - x_n^-}{x_n^+ - x_n^-} \in [0, 1] \tag{24}$$

All corners together form an enclosing hypercube around the off-grid state vector x^n, where each corner can be indexed by $c = (c_1, \ldots, c_d) \in \{0, 1\}^N$. $c_k = 0$ corresponds to the closest lower grid point, x_n^-, and $c_k = 1$ to the nearest higher grid point, x_n^+. Each point on the grid has an explicitly solved value, $V\big(x^{(c)}\big)$, computed within the iteration. Then, the interpolated value function $\tilde{V}(x)$ is then given by the weighted average of the 2^N corner values:

$$\tilde{V}(x) = \sum_{c \in \{0,1\}^d} w_c(t)\, V\Big(x^{(c)}\Big) \tag{25}$$

where the weight assigned to each corner is a function of its relative closeness to the off-grid state. Concretely, the weights of each corner are defined as the product of its coordinate-specific factors:

$$w_c(t) = \prod_{k=1}^{N} \begin{cases} 1 - t_n, & \text{if } c_k = 0 \\ t_n, & \text{if } c_k = 1 \end{cases} \tag{26}$$

where $x^{(c)}$ denotes the coordinates of the corner indexed by c. This approximation is linear within each cell, continuous across grid boundaries, and preserves corner values.

Figure 28 provides a graphic illustration for a 2x2-case. Let the grid size be 10, and the point to be evaluated $x = (2.5, 5)$. The surrounding grid points in this case are: (0,0), (10,0), (0,10), (10,10). Then, the point is 25% of the way between the lower and upper grid points for firm 1, and 50% of the way between the lower and upper grid points for firm 2. In the notation from above:

$$t_1 = 0.25$$
$$t_2 = 0.5$$

which gives the following weights for the surrounding grid points:

$$\begin{aligned} w_{00} &= (1 - t_1)(1 - t_2) = (1 - 0.25)(1 - 0.5) = 0.75 \times 0.5 = 0.375, \\ w_{10} &= t_1(1 - t_2) = 0.25 \times 0.5 = 0.125, \\ w_{01} &= (1 - t_1)t_2 = 0.75 \times 0.5 = 0.375, \\ w_{11} &= t_1 t_2 = 0.25 \times 0.5 = 0.125. \end{aligned}$$

The interpolated value is then obtained as:

$$\tilde{V}(2.5, 5) = 0.375\, V^{(0,0)} + 0.125\, V^{(10,0)} + 0.375\, V^{(0,10)} + 0.125\, V^{(10,10)} \tag{27}$$

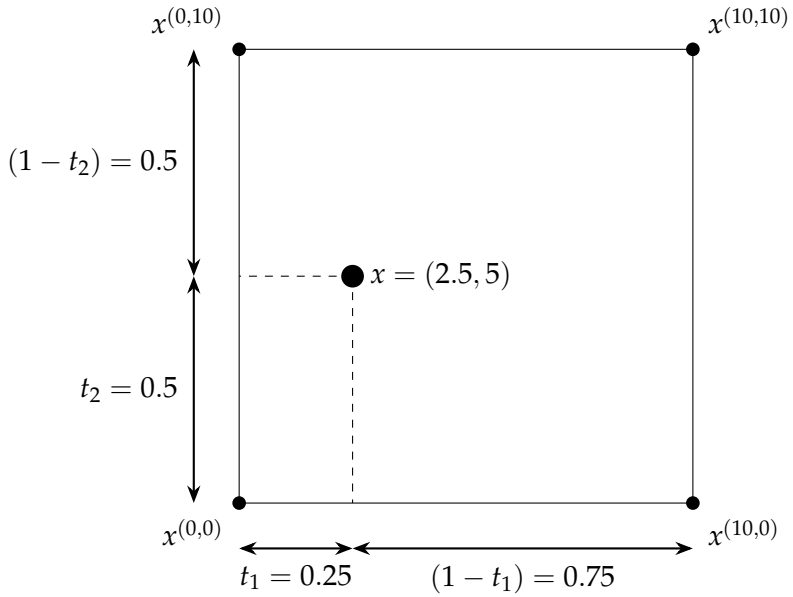


Figure 28: Illustration for the 2×2 case with $(t_1, t_2) = (0.25, 0.5)$.

Convergence and Numerical Details. The iterated solver uses a three-phase dampening schedule for the Bellman updates. In the first phase, a dampener of 0.8 is applied so that each update mixes 80% of the new Bellman value and 20% of the previous value, reducing overshooting early in the iteration. At iteration 20, the dampening factor is increased to 0.9 to accelerate convergence to the fixed point. Once the iteration count exceeds 245 iterations without convergence, both the value and policy dampeners are set to 0.5 to smooth out remaining numerical fluctuations and help the algorithm converge. Convergence is evaluated on the value and policy, as a percentage change of the value of each state, and the change in probability mass in each CCP. The convergence thresholds are set to 1e-4. When maxIter is reached without convergence, the solver stops and accepts the last dampened iterate; it flags a potential convergence issue and outputs the results in a CSV, so that convergence issues can be diagnosed and addressed. For the paper's estimated parameters and specifications, this is not the case. It can, for example, occur in some robustness checks with different cost parameters, most notably when the private shock variance used is insufficient.