

CLIMATE POLICY, INDUSTRY CONCENTRATION AND TECHNOLOGY ADOPTION: TAXES, SUBSIDIES AND REBATES*

Niklas Schoch

This version: October 12, 2026

Job Market Paper

[Latest version](#)

Abstract

Decarbonising heavy industry is a key challenge for climate policy. To reduce emissions, the European Union explicitly identifies the adoption of green technologies that lower emission intensity as a major avenue. This paper investigates which policies can incentivise an efficient decarbonisation of the heavy industry sector. I develop a dynamic model of imperfect competition and endogenous production technology upgrading, using recent developments in the modelling of dynamic continuous investment choices. A carbon tax alone leads to inefficient outcomes and a suboptimally low uptake of green technologies. The reason is that a carbon price low enough to avoid excessive output contraction and loss of market share to imports does not trigger sufficient investment in green technologies. In contrast, combining a tax with either an adoption subsidy or an output-based rebate increases domestic welfare. To quantify these effects, I estimate the model for the French cement industry, using a novel dataset that combines administrative data with hand-collected firm announcements on the low-carbon investments following the recent escalation of the EU Emissions Trading System carbon price. Compared to the optimal carbon tax alone, a suitable combination of a tax and an adoption subsidy increases domestic welfare by up to 4.8% (and reduces emissions by 5.1%), and an output-based rebate tied to an efficiency benchmark can yield a domestic welfare gain of up to 8.5% (and reduces emissions by 14.9%). Once a carbon border adjustment extends the carbon price to imports, the gain from a subsidy largely vanishes, and that from an output-based rebate falls to 2.1%.

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

JEL Codes: Q58, L13, H23, Q52, L61, C73

*University of Geneva, Graduate School of Economics and Management; niklas.schoch@unige.ch.

I am indebted to Johannes Boehm and Stefan Pollinger for their incredible supervision and support. I am also deeply grateful to Gautam Gowrisankaran for invaluable guidance, especially considering the technical aspects of this paper. Further, I want to thank Isabelle Méjean, who, in addition to countless helpful comments, provided me with access to the data, vital for this project. I would like to thank Haoge Chang, Nicolas Cœurdacier, Emma Du Puy, Antoine Ferey, Junnan He, Marleen Marra, Thierry Mayer, Kevin O'Rourke, Franz Ostrizek, Noémie Pinardon-Touati, Mar Reguant, Andrey Simonov, and seminar participants for their helpful comments. The computing resources for this project were provided by the Nuvolos Fellowship Programme 2026.

1 Introduction

Industrial emissions pose a significant obstacle to decarbonising global economic activity. In the EU, emissions from industrial installations accounted for more than 20% of total greenhouse gas emissions in 2025 (Eurostat, 2026). The European Union has the goal of reaching climate neutrality by 2050, and specifically identifies the adoption of green technologies that lower emission intensity as a major avenue to achieve emission reductions (EP, 2023). Economists widely recommend a carbon tax as the preferred policy option to address an externality such as climate change, since in a competitive industry, it can correct the externality and achieve the first-best outcome. Accordingly, carbon pricing, in the form of a tax or a permit market, is typically a central element of climate policy strategies. The EU, for example, has introduced a carbon price in the form of an emissions trading system (EU ETS). However, the most polluting industries targeted by climate policies, most notably cement, but also steel and chemicals, are typically highly concentrated (Koltay & Lőrincz, 2021) and compete with imports that are not subject to the same carbon price.

In this paper, I show that in such industries, carbon pricing alone is inefficient and leads to sub-optimally low uptake of green technologies, thereby falling short of one of the stated objectives of EU climate policy. The reason is that a carbon price low enough to limit the output contraction caused by market power and the loss of market share to unpriced imports does not induce sufficient low-carbon investment incentives. I evaluate several complementary policy options, reflecting current EU policies in addition to the EU ETS. When imports are not subject to the carbon price, both a green adoption subsidy and an output-based rebate, similar to the free allocation of allowances under the EU ETS, introduced in addition to a carbon tax, substantially improve domestic welfare. A carbon border adjustment mechanism (CBAM), extending the carbon price to imports, can partially substitute the other two instruments. The gains from a subsidy largely vanish, whereas the output-based rebate continues to raise domestic welfare.

To quantitatively evaluate decarbonisation policies and green technology upgrading in heavy industries, I develop a model of imperfect competition and strategic abatement technology choice. Firms compete in regionally segmented Cournot markets under capacity constraints and face a cost for their emissions due to environmental regulations (Ryan, 2012; Fowlie et al., 2016). The firms also face import competition, which is not subject to the domestic carbon price unless a CBAM is introduced. Over time, firms can invest in green technologies that reduce the emission intensity of their output, thereby altering their competitive position in future periods. Specifically, firms choose how much to reduce their emission intensity relative to a baseline technology, rather than making a binary adoption decision. This reflects a key feature of the data: all firms invest, and the variation lies in the extent of their investment. Abatement becomes increasingly costly as firms further reduce their emission intensity, and firms face stochastic shocks to the marginal cost of investment, generating heterogeneity in investment behaviour across firms and periods. In this setting, investment decisions are strategic. By lowering its emission intensity, a firm reduces its marginal cost under the carbon price and gains market share at the expense of its competitors. I model this dynamic game using the framework for continuous investment choices of Gowrisankaran and Schmidt-Dengler (2024). To compare policies, I compute domestic welfare, referred to simply as welfare throughout the paper, as the sum of consumer surplus, domestic profits and carbon revenue, net of investment costs and emission damages, including the damages from emissions embedded in imports.




Figure 1: EU ETS Price and Green Investment in the French Cement Sector

Notes: Carbon prices from ICAP Allowance Price Explorer, investment data from INSEE ANTIPOL, and hand-collected firm announcements.

I apply the model to the French cement industry, a highly concentrated and emission-intensive sector regulated under the EU ETS, which has seen substantial green investments in recent years. To estimate the static part of the model, I combine regional demand and import data from the industry organisation France Cement with plant-level capacities and production volumes inferred from verified emissions under the EU ETS. Following Ryan (2012), I estimate the production cost parameters by matching observed production quantities and firms' first-order conditions. To estimate the dynamic abatement cost parameters, I construct a novel dataset on green investments in the French cement industry, which combines administrative data with hand-collected firm announcements of green investment projects. Figure 1 shows that a wave of low-carbon investments has accompanied the recent escalation of the carbon price within the EU ETS, generating the variation I use to quantify the cost of adopting green technology for French cement firms. I employ a two-step, nested fixed-point GMM estimator. I first estimate competitors' investment policies from the firms' observed investment behaviour, and then let individual firms optimise against this response rule, matching the model-implied investment choices to firms' observed investment responses to the rising carbon price under the EU ETS.

In counterfactual simulations, I find that, absent a border adjustment, the domestic welfare-maximising carbon tax alone lies between the Pigouvian level and the optimal tax implied by the static analysis proposed by Buchanan (1969). For a social cost of carbon of €150/tonne, I estimate an optimal carbon tax of €135/tonne. By contrast, the statically optimal tax, which ignores the investment margin, is €90/tonne. This gap illustrates the tension between the output and the investment margin. While the lower tax limits output contraction and leakage, it provides insufficient incentives to invest in green technology. Setting the tax at the statically optimal level of €90/tonne results in a domestic welfare loss of 3.4% relative to the optimal tax of €135/tonne. Given the added investment incentives, the optimal tax reduces emissions by a further 37.7% compared with the statically optimal tax, while reducing consumer surplus by only 6.1%.

Adding a green investment subsidy to the carbon tax to counteract suboptimally low invest-

ment, I estimate that the optimal combination is an explicit investment cost subsidy of 60% with a tax rate of €70/tonne. Since the subsidy strengthens investment incentives directly, the tax can be set substantially lower, limiting output contraction and leakage. This combination increases domestic welfare by 4.8% relative to the optimal tax alone, while reducing emissions by a further 5.1%. Under an output-based rebate, firms receive a rebate based on their output and a fixed emission-intensity benchmark, independently of their actual emissions, a design that closely resembles the free allocation of allowances currently granted to the cement industry under the EU ETS. Using a benchmark of 0.2 kg of CO₂ per kg of cement, I find an optimal carbon tax of €165/tonne, yielding a domestic welfare gain of 8.5% and a further emission reduction of 14.9% relative to the optimal tax alone (€135/tonne). By supporting domestic production, the rebate allows a higher carbon price that strengthens investment incentives without excessive output contraction. It outperforms the subsidy since it acts as a production subsidy even for the cleanest firms. Neither of these two policies is self-financing: carbon revenues cover less than a third of subsidy payments, and the rebate yields negative carbon revenues. The domestic welfare gains over the optimal tax alone persist up to a cost of public funds of 13% for the subsidy and 20% for the rebate.

With a CBAM, the optimal tax rises to €175/tonne, above the social cost of carbon, and emissions fall by a further 52.2% relative to the optimal tax alone. Since importers now pay for their embedded emissions, a higher carbon price no longer shifts production and rents abroad, and domestic producers can pass the carbon cost on without losing market share, so the tax itself provides stronger investment incentives. Alongside the CBAM, the additional domestic welfare gain of an output-based rebate falls from 8.5% to 2.1%, and that of an adoption subsidy largely vanishes. Moreover, the rebate now slightly raises emissions by 5.9% relative to the CBAM alone. The domestic welfare gain of the output-based rebate still comes from subsidising the production of even the cleanest firms. These results suggest that the planned phase-out of free allocation forgoes part of the domestic welfare gains from an output-based rebate, but that the introduction of the CBAM, which protects domestic producers against unregulated imports more directly, substantially limits these losses.

This paper extends the literature on carbon pricing in the presence of market power by explicitly modelling the interactions with endogenous abatement investments. It has long been recognised that market power interacts with the effects of taxing an externality (Buchanan, 1969). More recently, studies have highlighted the implications for climate policy (e.g., Mansur 2007; Ryan 2012; Fowlie et al. 2016; Leslie 2018; Preonas 2024). Investments to reduce the emission intensity of production are the main way to achieve decarbonisation without halting domestic industrial production. This issue is thus becoming increasingly important amid growing backlash over deindustrialisation driven by climate policy. In addition, as suggested in Figure 1 and confirmed by causally identified reduced-form studies (Fowlie, 2010; Colmer et al., 2025), firms indeed respond to regulations by making capital investments. This paper quantitatively explores the relevance of this investment margin for the design of carbon policies. I find that the gains from carbon policies that address firms' incentives to invest in green technologies can be substantial, as a carbon tax alone cannot create these incentives optimally, especially when imports are not subject to a carbon price.

Studying this issue, I contribute to the structural IO literature on the implications of environmental regulations (e.g., Ryan 2012; Fowlie et al. 2016; Ganapati et al. 2020; Gowrisankaran et al. 2024; Toyama 2024; Gowrisankaran et al. 2025). To my knowledge, the endogenous choice of abatement investment in a structural model has so far been studied in settings without strategic interaction

between firms, such as price-taking or regulated electricity generators (Toyama, 2024; Gowrisankaran et al., 2025). I advance the literature by introducing the investment margin into a dynamic oligopoly setting. In particular, I use the framework for continuous investment choices of Gowrisankaran and Schmidt-Dengler (2024) to estimate and solve a multi-agent dynamic game of abatement investment, explicitly reflecting the prevalent market concentration in heavy industries.

Furthermore, my results complement the literature employing econometric techniques to estimate observed effects of carbon pricing (e.g., Fabra and Reguant 2014; Martin et al. 2014; Calel and Dechezleprêtre 2016; Wagner et al. 2016; Calel 2020; Dechezleprêtre et al. 2023; Colmer et al. 2025). Most of these empirical studies find little evidence of output contractions or other adverse effects of climate policies (e.g., 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 targeted policy implications to be derived.

Finally, this paper relates to the early theoretical literature on technology investments and environmental regulations (e.g., Delbono and Denicolo 1991; Laffont and Tirole 1996; Yi 1999; Montero 2002; Requate and 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 technology in a continuous choice setting, as well as the domestic welfare implications of a wide range of alternative policies.

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

2 Institutional Setting

This paper studies the French hydraulic cement market. Mixed with sand, water, and other aggregates, hydraulic cement forms concrete, a material used extensively in all construction sectors. Hydraulic cement is produced by heating limestone to high temperatures. Figure 2 illustrates a typical cement production process. This involves feeding limestone into a rotary kiln. After reaching over 1400°C, the limestone turns into cement clinker, which is ground up and mixed with additives to produce cement. Within the EU, approximately 90% of cement is sold in bulk to construction firms. The remaining 10% is sold directly to end consumers (CEMBUREAU, 2025).¹

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. The emission intensity of cement production, absent dedicated low-emission investments, is approximately 0.63 tonnes of CO₂-equivalent per ton of cement produced (ADEME, 2023). Emissions arise

¹The 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.




Figure 2: Cement Production Line (Marie et al., 2018)

from two sources: chemical processes that occur when limestone is heated (process emissions) and heat generation, 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 as a main abatement lever (Fowlie et al., 2016). However, given its considerable efficiency gains, the French cement industry had already broadly adopted dry production, the most efficient commercially employed technology, before the introduction of the EU ETS.² Starting from this technology standard, I observe three abatement levers being employed. Firstly, the fossil fuel used to heat the kiln can be replaced by biofuel feedstock to reduce energy emissions. This requires retrofitting the kiln's firing system. Secondly, a newly developed technology allows replacing parts of the clinker with calcined clay, which emits less CO₂ when heated. Producing calcined clay cement requires significant upgrades to the rotary kiln. Thirdly, remaining emissions can be captured and stored underground (CCUS).

Figure 3 plots emission savings by technology type, showing that all levers play a substantial role in firms' investment strategies.³ 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 investment cost increases with the desired biofuel tolerance. Similarly, the maximum replacement rate of clinker with clay is not fixed and can vary 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 other abatement levers reduce the CO₂ concentration in the flue gas stack. The individual levers can also 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, Holcim, EQIOM, and Ciments Calcia - accounting for over 90% of production capacity. To-

²The remaining wet kiln retrofits within my sample period are included in the estimation alongside the other low-carbon investments.

³"Integrated" refers to investment events in which I observe kiln upgrades deploying at least two of the discussed technology options simultaneously.




Figure 3: Share of Emissions Abated by Technology

Notes: The shares plotted are based on data on low-carbon investments from INSEE (ANTIPOL) and hand-collected firm announcements, for the years 2012-2025. For details on the data, please refer to Section 4.

gether, these four firms operate 25 integrated cement plants throughout France.⁴ The HHI values range from 0.4 to 0.6, indicating significant market concentration across all markets during the sample period. In my sample period, the aggregate installed capacity of integrated cement plants within France amounts to roughly 19 Mt/year.⁵ Generally, competition in the cement industry is regionally segmented due to high transportation cost relative to cement's 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 150 km of the production facility, and 90% within 250 km (EC, 2014). Although high inland transport costs limit trade within France, domestic producers face import competition by sea and from neighbouring countries. The import share in France hovers slightly above 10% on average, with individual regions reaching up to 20%.

In line with the available demand data, I follow the industry organisation France Ciment and divide the French market into four regions, which I treat as independent markets. To determine which plants serve which market, I allocate plants based on their location.⁶ Figure 4 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, treating firms within each market as a single agent. Their capacity constraint is derived from the sum of the capacities of the plants of a given firm in this market.

The EU ETS. The European Union's cement industry has been covered by the EU ETS since 2005. Importantly, French cement installations represent a small share of covered emissions and are thus

⁴Integrated cement plants include both clinker production and grinding, as opposed to facilities only grinding ready clinker into cement powder.

⁵There have been changes in capacity in my sample period, mostly retirements of older plants. I treat the occurring capacity changes as exogenous, but anticipated by all agents. The details are discussed in Section 4.

⁶Robustness checks in which borderline cases are assigned differently mostly return comparable estimates.




Figure 4: Regions and Plant Locations

Notes: The regions plotted depict the division of France into the four regions for which disaggregated demand data are available. Shapes of the plants indicate the owner, and the colour matches the market to which the plant is allocated.




Figure 5: Share of Carbon Price in Marginal Cost

Notes: Based on the estimated production costs, excluding compliance cost, of €72/ton, and the EU ETS price depicted in Figure 1. The final data point reflects the assumed long-term carbon price expectation of €125/ton from 2026 onwards, discussed in Section 5. The cost estimation is reported in Section 5.

price-takers, not influencing the carbon price itself. As shown in Figure 1, a 2018 reform of the EU ETS, combined with the post-COVID economic recovery, has driven the carbon price up significantly, now hovering around €70-€100/ton since 2022. Given cement's high emission intensity, these carbon price levels are a major cost driver for cement production. Figure 5 plots the share of total marginal cost represented by the EU ETS carbon price over time, at a fixed emission intensity of 0.63, the industry average prior to any meaningful abatement investments (ADEME, 2023).

In 2025, the average EU ETS carbon price was around €75/ton. At full emission intensity, around 40% of the marginal costs for a cement plant are 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 6 plots the realised carbon price burdens (dashed lines) against the theoretical burden had the firms not invested (black solid line, as in Figure 5). It shows that investments have translated into substantial marginal cost savings, with significant heterogeneity across firms. The heterogeneity in investment behaviour across firms and markets, and over time, constitutes the key variation exploited in the dynamic estimation of this paper.

Historically, cement installations receive a large share of their allowances for free, based on a clinker benchmark and historical production, independently of their own emission intensity. I discuss how I account for free allocation and carbon price expectations in Section 5. In addition to the carbon price increase, investment 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 with the carbon price regulation in the dynamic estimation.




Figure 6: Share of Carbon Price in Marginal Cost with and without Investment

Notes: Based on the estimated production cost, reported in Section 5, the EU ETS price depicted in Figure 1, and hand-collected announcements of firms' low-carbon investment projects.

3 Model

The model captures the characteristics of heavy industries that drive my results: firms compete imperfectly, face an import fringe, and they can invest to reduce the emission intensity of their output. Accordingly, the model contains two parts: a static Cournot competition in every period, and a dynamic game of abatement investment linking the periods.

In the static part, firms compete in each regional 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. This part largely follows the established setup by Ryan (2012) and Fowlie et al. (2016). In the dynamic part, firms can invest to reduce the emission intensity of their output. Since emission intensity enters the marginal cost through the carbon price, abatement investment alters a firm's competitive position in all subsequent periods.

The dynamic part is where my model departs from existing structural models of environmental regulation. Previous studies focus on which firm invests, modelling a stochastic shock to a fixed investment cost, with abatement primarily achieved through the energy efficiency of new entrants (Ryan, 2012; Fowlie et al., 2016). In my setting, all firms invest, and the variation lies in the extent of their investment. Reflecting the menu of scalable and combinable abatement levers available to firms (see Section 2), I therefore model abatement as a continuous investment choice of incumbent firms. Following Gowrisankaran and Schmidt-Dengler (2024), firms receive idiosyncratic shocks to the marginal cost of investment. This allows me to compute the equilibrium of a dynamic oligopoly

game with continuous abatement investment, and to estimate the abatement investment cost function from firms' observed investment responses.

The policies I study enter the model at different points. A carbon tax and an output-based rebate alter firms' compliance costs in the static part, whereas a technology adoption subsidy reduces investment costs in the dynamic part. A CBAM affects the import supply curve. For each policy, I compute the resulting equilibrium and evaluate welfare as the sum of consumer surplus, firm profits and carbon revenue, minus investment costs and emission damages, including leakage through imports.

I set up an infinite-horizon model in discrete time. 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 at the end of the period, affecting their intensity only in the following period. The state transitions are deterministic given the investment decisions. The remainder of this section describes each model element in detail.

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

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

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

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

where $I_{m,t}$ denotes the aggregate imports into market m in period t , $\rho_{m,t}^0$ 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.⁷ 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 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 starts to escalate. This cost structure reflects, for example, the extra cost incurred by forgoing maintenance, which requires an annual downtime in production, or overtime pay for night-time and weekend shifts. The production costs are:

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

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.

⁷A detailed list of the fringe firms, and their date of entry into the French market, is provided in Appendix A, Table 9.

Environmental Compliance Cost. $\varphi(q_i, \eta_i, \tau)$ denotes the environmental compliance costs, which are part of the production cost. A firm's compliance costs are determined by its taxable total emissions, its output q_i multiplied by its emission intensity η_i , and the carbon price τ . The functional form of φ will vary by counterfactual policy. In the case of a carbon tax, it reads as:

$$\varphi(q_i, \eta_i, \tau) = \tau q_i \eta_i \quad (4)$$

Under an output-based rebate, firms receive a rebate based on their output and a defined efficiency benchmark. The compliance cost, $\varphi(q_i, \eta_i, \tau)$, for firm i in this case, is then:

$$\varphi(q_i, \eta_i, \tau) = \tau q_i (\eta_i - \eta^*) \quad (5)$$

where η^* is a defined efficiency benchmark, set by the policy maker.

Under a CBAM, importers pay the domestic carbon price on the emissions embedded in their products, evaluated at the import emission intensity $\eta_{Imports}$. I model this by letting the import supply curve respond to the market price, minus the carbon cost they must pay (Fowlie et al., 2016).

$$\ln I_{m,t}(P_{m,t}; \rho) = \rho_{m,t}^0 + \rho^1 \ln(P_{m,t} - \tau \eta_{Imports}) \quad (6)$$

The import emission intensity is fixed, following Fowlie et al. (2016), and importers respond to the CBAM only through quantity adjustments.

Firm Profits. Given production and compliance costs, firms choose production quantities to maximise per-period profit. The corresponding optimisation problem is given by:

$$\max_{q_i} \Pi_i(q_i; \delta, \eta, \tau) = P(Q)q_i - C(q_i; \delta, \eta_i, \tau) \quad (7)$$

In the following, $\Pi_i^*(\eta; \delta, \tau)$ will denote the optimised, equilibrium per-period profit, given δ, η, τ and the functional form of φ , determined by the policy.

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 2), I model firms as facing a continuous abatement investment cost function:

$$InvCost_i(\theta_i, \lambda, \varepsilon_i) = s_i \theta_i (\lambda_1 + \lambda_2 \theta_i + \lambda_3 \varepsilon_i) \quad (8)$$

where $\theta_i \in [0, 100]$ represents the percentage reduction in emission intensity compared to a baseline intensity $\bar{\eta}$, denoting the intensity absent any abatement efforts.⁸ I set the baseline intensity to 0.63 tons of CO₂-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 :

$$\eta_i = \frac{100 - \theta_i}{100} \bar{\eta} \quad (9)$$

I will use these terms interchangeably. Following the framework by Gowrisankaran and Schmidt-Dengler (2024), I introduce firm-level heterogeneity in the form of cost shocks affecting the marginal cost of investment, $\varepsilon \sim N(0, 1)$, scaled by λ_3 . The model allows for matching the observed variation in

⁸The model implies the possibility of negative investment cost for small investments. These can be thought of as reflecting natural efficiency gains in the absence of explicit abatement efforts.

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 investment cost 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:

$$InvCost(\theta'_i|\theta_i) = InvCost_i(\theta'_i, \lambda, \varepsilon_i) - InvCost_i(\theta_i, \lambda, \varepsilon_i) \quad (10)$$

With a technology adoption subsidy, firms pay only a fraction of the investment cost. For the firm, the cost reads as:

$$\text{Realised InvCost} = (1 - \zeta)InvCost(\theta'_i|\theta_i) \quad (11)$$

where ζ is the subsidy share. In the welfare calculation, I assume public funds cover the remaining cost.

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:

$$V_i(\theta; \delta, \lambda, \varepsilon_i, \sigma(\theta)) = \max_{\theta'_i} \Pi_i^*(\eta; \delta, \tau) - InvCost(\theta'_i|\theta_i) + \beta \int E_{\varepsilon_i} [V_i(\theta'; \delta, \lambda)] dP(\theta'_{-i}, \sigma(\theta)) \quad (12)$$

where θ denotes the vector of investment levels of all firms, $P(\theta'_{-i}, \sigma(\theta))$ is the probability distribution of competitor states, a function of the starting state vector and their strategies σ . β denotes the discount factor. The results in this paper were computed with $\beta = 0.975$.

Welfare. In the model, the welfare of the domestic economy is the sum of consumer surplus, firm profits and carbon revenue, minus investment costs and damages from emissions.

$$W = CS(Q^{Dom} + Imports) + \sum_i Profits + \text{Carbon Revenue} - \sum_i InvCost - \text{Damages}(E^{Dom} + Leakage) \quad (13)$$

Carbon revenues are paid by firms and rebated in a lump sum. Thus, the collection of carbon revenues is itself a welfare-neutral process in my setup. I assume an emission damage function with constant marginal damages:

$$\text{Damages}(E) = \gamma E \quad (14)$$

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

$$E = E^{Dom} + Leakage = \sum_{j=1}^N \eta_j q_j + I\eta_{Imports} \quad (15)$$

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 ($\eta_{Imports} = 0.63$). CBAM revenues are paid by the import fringe, whose profits are not included in domestic welfare. The collected CBAM revenues are thus counted towards domestic welfare. This is discussed in more detail when analysing the outcomes under a CBAM in section 6.

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.⁹ 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:

$$-\Delta C = -\tau\Delta\eta_i q_i \quad (16)$$

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:

$$-\Delta MC = -\tau\Delta\eta_i \quad (17)$$

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)):

$$-\frac{\partial \Pi_i^*}{\partial \eta_i} = q_i \left[\underbrace{\tau}_{\text{Direct saving}} - \underbrace{P'(Q) \frac{d(Q - q_i)}{d\eta_i}}_{\text{Price hike due to comp. contraction}} \right] \quad (18)$$

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

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 set the cost of achieving this reduction against their private profit, while 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

⁹The mechanisms are equivalent for any regulation where the compliance cost increases in emissions.

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. As in the previous paragraph, I illustrate the mechanisms for the case of a carbon tax, where carbon revenues are $R = \tau E^{Dom}$. Using the welfare definition in (13), the welfare effect of a marginal intensity reduction, to be compared to the investment cost, can be decomposed as

$$\Delta W = \sum_{j=1}^N \Delta \Pi_j + \Delta CS + \Delta R - \gamma \Delta E = \sum_{j=1}^N \Delta \Pi_j + \Delta CS - (\gamma - \tau) \Delta E^{Dom} - \gamma \Delta Leakage \quad (19)$$

where $\Delta \Pi_j$ is the change in firm j 's optimised per-period profit from (7), net of compliance costs, ΔCS is the change in consumer surplus, and γ is the constant marginal damage parameter. Since firms already pay τ per tonne of domestic emissions, only the difference between the social cost of carbon and the carbon tax remains as an externality on domestic emissions, while leakage through imports is not priced.

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 remaining externality. Rearranging (19) yields

$$\Delta W - \Delta \Pi_i = \underbrace{\sum_{j \neq i} \Delta \Pi_j}_{<0 \text{ (Business stealing)}} + \underbrace{\Delta CS}_{>0 \text{ (Market expansion)}} - \underbrace{(\gamma - \tau) \Delta E^{Dom}}_{\geq 0 \text{ if } \tau \leq \gamma \text{ (Residual externality)}} - \underbrace{\gamma \Delta Leakage}_{>0 \text{ (Leakage reduction)}} \quad (20)$$

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. Hence, part of the private gain $\Delta \Pi_i$ reflects a redistribution from rivals rather than net surplus creation. Even though firm i does not consider the change in consumer surplus and the remaining externality, its unilateral private incentives can therefore exceed the welfare gain through the business-stealing channel. 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.¹⁰

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 from the comparison with aggregate profits, and the welfare change is

$$\Delta W = \Delta \Pi + \Delta CS - (\gamma - \tau) \Delta E^{Dom} - \gamma \Delta Leakage, \quad \Delta \Pi \equiv \sum_{j=1}^N \Delta \Pi_j. \quad (21)$$

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

$$\Delta W - \Delta \Pi = \underbrace{\Delta CS}_{>0 \text{ (Market expansion)}} - \underbrace{(\gamma - \tau) \Delta E^{Dom}}_{\geq 0 \text{ if } \tau \leq \gamma \text{ (Residual externality)}} - \underbrace{\gamma \Delta Leakage}_{>0 \text{ (Leakage reduction)}}. \quad (22)$$

Since firms do not internalise the gains in consumer surplus and the reduction in leakage, and the carbon tax internalises the domestic externality only up to τ , common private incentives fall short

¹⁰I abstract from the rare, theoretical case where a reduction in emission intensity by one firm can increase aggregate emissions.

of the social incentives for any $\tau \leq \gamma$. Even at $\tau = \gamma$, the wedge reduces to the gains in consumer surplus and leakage reduction, which remain positive. Whether private incentives fall short of social incentives in equilibrium therefore depends crucially on the expectations and strategies of all firms. My full model will account for these mechanisms, in which each firm's gain is derived from expected profit gains, given expectations of other firms' strategies.

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 a stochastic optimisation procedure, Metropolis-within-Gibbs sampling in the spirit of Chernozhukov and Hong (2003), using the implementation of Ryan (2012), to find the production quantities that minimise the gap 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 and price pair that clears the market, using the market-level demand and import supply functions. 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 and depend on the current state vector. The equilibrium strategies, $\sigma_i^*(\theta)$, mapping states into choice probabilities for investment levels, $\sigma : \theta \rightarrow P(\theta')$, satisfy:

$$V_i(\theta; \varepsilon_i, \sigma_i^*(\theta), \sigma_{-i}(\theta)) \geq V_i(\theta; \varepsilon_i, \sigma'_i(\theta), \sigma_{-i}(\theta)) \quad \forall i, \sigma'_i(\theta) \quad (23)$$

The results in this version of the paper were computed on a state space with a grid step of size 20%.¹¹ Thus, I explicitly compute the per-period payoff and option value for all possible combinations of firms' emission intensities in 20% increments.

Each firm's investment choice is represented as a discretised set of available investment bins (Gowrisankaran & Schmidt-Dengler, 2024). I allow firms to choose their optimal emission-intensity reductions in increments of 5%. Each bin corresponds to all investment levels between two allowed investment increments, e.g., in my case, (0, 5], (5, 10]. 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 that depend on the relative position of the off-grid point along each dimension. Details of the interpolation algorithm are provided in Appendix C.

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 defined interval, 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

¹¹The North West market with four players is an exception, where I set the grid to 25% to manage computational cost. Experimentation with finer grids yielded the same results.

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. 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 compute the dynamic equilibrium by embedding this algorithm into value-function iterations, solved using iterated best responses. 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 at a later point, 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 the competitors 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 and industry paths, but rather a vector of choice probabilities for each situation in the game that firms can find themselves in.¹²

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 (Fowlie et al., 2016). Letting the industry start at full emission intensity of 0.63 (see ADEME (2023), discussed in Section 2), 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).

4 Data

In this section, I discuss the various data sources I use to estimate the model and simulate the policy scenarios. Table 5 in Appendix A provides descriptive statistics for the main variables, and more detailed information on the data sources is given in Table 6.

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 4). For prices, I use administrative data on firms' self-reported prices from INSEE (EAP). For each market and year, I compute the market price as the average of the firm-level prices from all firms active in that market, weighted by production level. Given the high transportation costs, imports play a minor role in the French cement market, accounting for slightly above 10% of aggregate demand. Nonetheless, in counterfactuals with high carbon prices, the import margin can affect the quantitative results. Since

¹²The discretisation can lead to numerical issues. For some combinations of grid step size and model parameters, non-convergence can occur in 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.

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 17.

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.¹³ I allocate each firm's production facilities to one of the four markets by location. Figure 4 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 7 and 8.¹⁴ In 2015, Holcim and Lafarge merged. Using the European Commission's merger decision, I track the resulting ownership changes (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.

Production. Physical production data for 2012–2019 come from INSEE (EAP) but are available only 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 7 reports a histogram of the ratio of inferred over observed production per firm and year.

¹³In the cases where there is a discrepancy, I assume the capacity that is closest to the production volume.

¹⁴Grinding 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%.




Figure 7: Inferred versus Observed Production

Notes: This graph plots the firm-level production quantities per year, inferred from emissions data from the EUTL, divided by the firm-level observed production data from INSEE (EAP).

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 clinker imports.¹⁵ Figure 23 in Appendix B shows the plot by firm, indicating limited, but present heterogeneity across firms. This heterogeneity is 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. Because the model recovers economic costs, including opportunity, overhead and financing costs, from firms' choices, the estimation requires the physical outcome of each investment rather than its engineering cost. I therefore convert reported spending into reductions in emission intensity.

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 the spending, reported in monetary values in ANTIPOL, into reductions in emission intensity. Figure 8 depicts the relationship I exploit between euros spent per capacity and percentage reductions in emission intensity. Since the administrative data are currently

¹⁵Portland 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 8: Engineering Abatement Investment Cost Curve

Notes: This graph plots the linear regression used to translate monetary spending into a reduction in emission intensities. The data points are based on hand-collected firm announcements, which are reported in detail in Table 10.

only available up to 2022, I further complement these 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 10. 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.¹⁶

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 data on investment subsidies received by French cement firms and, for abatement investments, hand-collect the most recently announced subsidies.

5 Estimation: Method and Results

This section discusses the estimation strategies separately for the model's static and dynamic components. I first discuss the static estimation, since the results feed into the dynamic estimation by determining the static payoffs in the dynamic optimisation. All objective functions in the estimations are minimised using the Covariance Matrix Adaptation Evolution Strategy (CMA-ES).¹⁷

¹⁶I use 2019 emissions as the reference year to avoid COVID-related distortions.

¹⁷For 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).

5.1 Static Parameters

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

$$\alpha^1 = -2 \quad (24)$$

$$\rho^1 = 0.5 \quad (25)$$

Using data on demand and prices, in combination with these elasticities, I fit year- and market-specific intercepts to construct demand and import supply curves for each 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 between observed production quantities and the model predictions and (2) the squared difference between marginal revenue and marginal cost at the observed quantities.¹⁸ 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 that year. This lets me isolate the constant part of the marginal cost curve, net of EU ETS compliance costs. Secondly, free allowances in Phase III (2013-2020) were allocated based on historical production levels, but reduced if firms produced less than 50% of their historic production volume (see Branger et al., 2015, for details). This can, in principle, create incentives to alter production decisions. In my sample, output remains well above 50% of historical levels, so the threshold is never close to binding, and free allowances are unlikely to affect the choice of production level.

The results of the estimation are reported in Table 1.

Table 1: Estimated Cost Function Parameters

Parameter	Estimate
MC (δ_1)	71.74 EUR/ton
Utilisation threshold (δ_2)	76.66%
Utilisation cost (δ_3)	50.5

Notes: Parameters refer to the production cost function, Equation 3.

The utilisation cost implies that, after exceeding 77% of capacity utilisation, the marginal cost of firms increases by €1 with every additional percentage point of utilisation. The estimated constant marginal cost is €72/ton, and the average price for a ton of cement during my sample period is €112/ton. Excluding EU ETS compliance costs, this implies an operating margin ranging between 30% and 40% across markets and years. I plot the profit margins implied by my estimates per firm and

¹⁸The estimator has equally been employed by Fowlie et al. (2016).




Figure 9: Carbon Pricing Regimes and Investment Periods

Notes: Carbon prices from ICAP Allowance Price Explorer, Investment data from INSEE ANTIPOL, and hand-collected firm announcements.

market over time in Appendix B, Figures 19 and 18. 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 20, 21, and 22.

5.2 Dynamic Parameters

As is common in the literature (Aguirregabiria et al., 2021), I employ a two-step, nested fixed-point GMM estimator to estimate the parameters governing the dynamic abatement investment decision. Estimating the dynamic game directly would require solving for the equilibrium of the game for each parameter guess, which is computationally expensive. To avoid this computational burden, I first estimate competitors' investment policies from the firms' observed investment behaviour. Given the small number of firms in my sample, I cannot estimate these policies flexibly. Instead, I approximate them with a simple reduced-form rule in a few market observables. In the second step, I let individual firms optimise against this response rule, meaning they assume their competitors follow the estimated reduced-form policy. This turns the dynamic game into a single-agent optimisation for each firm. For each parameter guess, I then 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. This allows me to generate the model-implied investment behaviour for each period, which I can match against observed investment moments.

Identifying Variation and Investment Periods. The low-carbon investments studied here involve significant construction times. Given different carbon pricing regimes, I estimate the model using three-year investment periods, followed by a terminal period from 2026 onwards. Carbon prices over these periods are shown in Figure 9.




Figure 10: Subsidy Rate per Period

Notes: The rate reflects the total subsidies paid in each period, reported in the European State Aid Transparency Database, divided by the total observed investment amount.

For the periods after 2019, for which demand data are not available, I hold the demand and import supply curves fixed at their 2019 levels in each market, and compute the static payoffs using the observed capacities and ownership structure in each period. I assume that firms face a long-term price expectation of €125/ton, which is realised from 2026 onwards.¹⁹ Then, I solve for the infinite continuation value at €125/ton and backward-induct the optimal 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 increases in the EU ETS allowance price 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 10 plots the observed subsidy rates for each investment period.²⁰

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

¹⁹The robustness of the estimation for other carbon price expectations is reported in Appendix B, Table 11.

²⁰I only include subsidies received for specific abatement projects, not R&D funding.

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:

$$\theta_{i,(t+1)} = \beta_1 \text{MarketShare}_{i,t} + \sum_{k=0}^4 \gamma_k 1\{\text{Period}_t = k\} + \beta_2 \text{stateExAnte}_{i,t} + \epsilon_{i,t}. \quad (26)$$

Given the small sample (55 observations), most coefficients are imprecisely estimated, but they have the expected signs. 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. Catching up after competitors have invested, when aggregate output is already high, is therefore less valuable than investing early. 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.

Table 2: Competitor Policy Functions

<i>Dependent variable:</i>	
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

Notes: *p<0.1; **p<0.05; ***p<0.01. This table shows the results of estimating Equation 26.

Reflecting the low number of observations, I do not observe firms in all possible states of the game and cannot fully identify the conditional choice probabilities empirically. Thus, I employ a full-solution estimator. The firms' behaviour I set against the data is derived from a full solution of the value function over time, given the reduced-form competitor strategies. This contrasts with the estimator proposed by Bajari et al. (2007), which relies more heavily on 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 using a nested fixed-point GMM estimator. I minimise the difference between the data and the model for the following moments: (1) mean abatement investment, (2) mean squared investment, (3) the share of positive investments, (4) the share of large investments (emission-intensity reductions of at least 50%). I bootstrap the efficient weighting matrix from the variance-covariance matrix of the data moments, and minimise the resulting GMM criterion:

$$\hat{\lambda} = \arg \min_{\lambda} \bar{g}(\lambda)' \hat{W} \bar{g}(\lambda) \quad (27)$$

where $\bar{g}(\lambda)$ is the vector of moment deviations and $\hat{W}$ is the efficient weighting matrix. The results are reported in Table 3.

Table 3: Estimates of Abatement Cost Function Parameters

	λ_1	λ_2	λ_3
Estimate	0.067	0.180	2.221
Std. error	3.092	0.055	0.796

Notes: Standard errors are GMM standard errors under efficient weighting, discussed in Appendix B.

As a reference point, these estimates imply an expected cost of about €8.7bn to reduce the emission intensity of the entire French cement industry by 50%.²¹ For comparison, the French environmental agency ADEME estimates capital expenditures of €4.4bn for a reduction in emission intensity of slightly less than 50%, net of the projected decline in demand (Mari & Sourisseau, 2021). This figure excludes the cost of capital. Unlike such techno-economic assessments, my estimates capture the economic cost of adoption revealed by firms' investment decisions, including opportunity costs associated with project overhead, downtime, and adjustment.

6 Counterfactuals

To evaluate counterfactual policy scenarios, I solve for the dynamic game's equilibrium and simulate the expected industry path forward using Monte Carlo simulations, assuming that every firm starts at full emission intensity.²² In the following sections, I discuss the equilibrium results for different carbon tax levels (Section 6.1), a technology cost subsidy (Section 6.2), and an output-based rebate (Section 6.3), and their interactions with a CBAM (Section 6.4). I present a full, tabular breakdown of all outcome variables under the optimised policies in Table 4, at the end of this section.

Following Fowlie et al. (2016), the results are presented as net present values (NPV). 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, and emission damages are set at €150/tonne. The robustness to other values for the emission damages is reported in Table 13, in Appendix D.

²¹This evaluates the investment cost function at the mean of the cost shock and thus corresponds to the expected cost of an industry-wide mandate. It abstracts from firms timing their investments to favourable cost draws, which lowers realised costs.

²²The details are discussed in Section 3.

6.1 Carbon Tax

To understand the effects of varying carbon tax levels, I contrast (1) the full dynamic model (Full Model) with (2) a static oligopoly model without the investment margin, in which firms keep their baseline emission intensity (Static Oligopoly). The statically optimal tax is the tax that maximises welfare in the static oligopoly model. 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. I discuss the most relevant outcome variables, plotted in Figure 11. A comprehensive set of plots for all outcomes is available in Appendix D.




Figure 11: Carbon Tax: Outcomes

Notes: Full Model: dynamic model with abatement investment; Static Oligopoly: model without the investment margin. Net present values, aggregated across markets. The dashed vertical line marks the social cost of carbon (€150/tonne).

Consumer surplus, determined by the aggregate market quantity, is plotted in Panel 11a. 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. Intuitively, emission damages (Panel 11b) decline with increasing tax levels. As the carbon tax increases, the decrease is stronger in the dynamic model. Importantly, the emission damages are substantially lower even though, as described above, the production quantity in the

complete dynamic model remains substantially higher than in the static setting. This illustrates well how adopting green technology benefits economic welfare. By decreasing the production's emission intensity, more can be produced, boosting consumer surplus, without incurring excessive emission damages. However, firms incur higher investment costs to achieve this (Panel 11c).²³ Firms' investment increases with a higher tax burden, reflecting an escalating benefit from saving on compliance costs and an increased relevance of relative emission intensity to a firm's competitive positioning.

Since imports are not subject to the carbon price, a higher tax also shifts production towards imports. Emissions from imports rise with the tax (Panel 11d), although this emission leakage remains limited: at the optimal tax discussed below, imports rise by 15% relative to no carbon pricing, offsetting 3.1% of the domestic emission reduction. For domestic welfare, a further cost is that importers gain market share and rents, which are not part of domestic welfare.




Figure 12: Carbon Tax: Total Welfare

Notes: Full Model: dynamic model with abatement investment; Static Oligopoly: model without the investment margin. Net present values in euros, aggregated across markets. The dashed vertical line marks the social cost of carbon (€150/tonne).

Putting all these components together yields the aggregate domestic welfare implications of the carbon price levels, plotted in Figure 12. 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, and because of the presence of leakage concerns. This is true in both the static and dynamic contexts. 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 the social cost of carbon. The optimal tax, lying in the middle, illustrates these countervailing effects when introducing carbon taxation in a

²³The model structure implies the possibility of negative investment cost. 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.

concentrated industry that can invest to reduce its carbon emissions. Market power already restricts the output, lowering the welfare benefits of carbon taxation. Further, high carbon taxes induce leakage towards high-emission imports. 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.

Quantitatively, with a constant damage of emissions of €150/tonne, I estimate an optimal carbon tax of €135/tonne, compared to €90/tonne implied by a static oligopoly analysis. Evaluated in the full model, setting the tax at the statically optimal level would result in a welfare loss of 3.4%. Compared with the statically optimal tax, the optimal tax reduces emissions, including those from imports, by a further 37.7%, while reducing consumer surplus by only 6.1%. One major reason is that, due to the additional investment incentives, the emission intensity of domestic production falls to 0.168 tonnes of CO₂ per tonne of cement at the optimal tax, less than half of the 0.352 at the statically optimal tax.

Still, firms invest less than would maximise welfare. I compare realised investment levels with a welfare-maximising benchmark in Figure 13. To compute this benchmark, I solve the value function iteration and the Monte Carlo simulation exactly as in the oligopoly model, but using welfare, defined in Equation 13, as the per-period payoff instead of the firm's profit. Doing so, I generate a benchmark investment behaviour that would result if firms invested to maximise welfare rather than profits, while keeping product-market behaviour fixed. Whilst an increasing carbon tax increases investment, the plot shows that firms' investment levels remain suboptimal and, crucially, well below this benchmark for the optimal tax of €135/tonne, as firms do not account for the gains in consumer surplus from investing.




Figure 13: Carbon Tax: Investment Cost, Private vs. Welfare-Maximising

Notes: Full Model: equilibrium investment of profit-maximising firms; Social Invest: investment that maximises welfare, holding product-market behaviour fixed. Net present values in euros, aggregated across markets. The dashed vertical line marks the social cost of carbon (€150/tonne).

6.2 Carbon Tax with Investment Subsidy

A natural response to address both the output and the investment margin is the introduction of a second policy instrument. Given that reduced investment incentives are a primary driver of welfare losses, an investment subsidy alongside a carbon tax seems a promising and widely employed policy mix.

Figure 14 shows the total welfare for varying technology subsidy shares, combined with the jointly optimal carbon tax of €70/tonne. The optimal technology subsidy of 60% of investment cost yields a welfare gain of 4.8% relative to the welfare attained under the optimal simple tax of €135/tonne, while reducing emissions by a further 5.1%.




Figure 14: Carbon Tax with Subsidy: Total Welfare

Notes: Carbon tax fixed at €70/tonne; the horizontal axis shows the subsidy as a share of investment cost. Net present values in euros, aggregated across markets.

The subsidy shifts abatement from output contraction to cleaner production (Table 4). Since the subsidy strengthens investment incentives directly, the carbon tax can be set substantially lower, limiting the output contraction and the loss of market share to imports. Compared with no carbon pricing, domestic production falls by only 10%, compared with 50% under the optimal tax alone. At the same time, the emission intensity of domestic production falls to 0.096, compared with 0.168 under the tax alone. Relative to the optimal tax alone, consumer surplus rises by €15.2 billion and domestic profits by €4.8 billion, while damages from emissions embedded in imports fall by €1.3 billion. Against these gains stand higher investment costs of €15.1 billion, carbon revenues that are €3.6 billion lower, and slightly higher domestic emission damages, leaving a net welfare gain of 4.8%.

This can be an attractive policy mix, since tax revenues can be used to fund the subsidy. For example, in the European Union, the EU ETS revenues support subsidy schemes for low-carbon investment in industry. Figure 15 plots the net public funds, meaning carbon tax revenues minus paid subsidies, for increasing technology cost subsidies. At the welfare-optimal subsidy of 60%, the policy is not self-funded. Carbon revenues of €4.1 billion cover less than a third of the subsidy payments of

€14.2 billion. The policy thus requires the mobilisation of €10.1 billion of public funds, and its welfare gain over the individual tax remains up to a maximum cost of public funds of 13%.




Figure 15: Carbon Tax with Subsidy: Net Public Funds

Notes: Net public funds are carbon tax revenues minus subsidy payments. Carbon tax fixed at €70/tonne; the horizontal axis shows the subsidy as a share of investment cost. Net present values in euros, aggregated across markets.

However, even this combination of tax and subsidy leaves residual distortions. Firms still hold market power and excessively restrict their output. At full emission intensity, lowering the carbon tax below the social cost of carbon allows an implicit production subsidy. However, when firms invest to become close to emission-neutral, a tax, which can never be below zero, cannot achieve such an implicit production subsidy.

6.3 Carbon Tax with an Output-Based Rebate

A widely adopted policy to avoid adverse effects on production volumes is an output-based rebate (OBA), similar to the free allocation of allowances under the EU ETS. Under this policy, firms receive a rebate based on their production and a fixed emission-intensity benchmark, independently of their actual emissions. Crucially, firms also receive the full rebate when their emission intensity lies below the benchmark, so that the rebate exceeds their carbon bill and acts as an explicit production subsidy, directly counteracting the market power distortion. In addition, given that the rebate is based on a set efficiency benchmark, the effective static tax rate is decoupled from the dynamic investment incentive. Thereby, the policy-maker can set a higher effective tax rate for dynamic investment incentives, whilst avoiding excessive output contraction in the static equilibrium. Performing a grid search over the benchmark and tax levels, I find that a benchmark of $\eta^* = 0.2$ with a carbon tax of €165/tonne maximises welfare, yielding an 8.5% welfare gain over the best simple tax, while reducing emissions by a further 14.9%. As shown in Figure 16a, the optimal tax rate under this scheme is higher than under the tax alone.




Figure 16: Carbon Tax with OBA: Outcomes

Notes: Tax: carbon tax alone; OBA: carbon tax with an output-based rebate at a benchmark of $\eta^* = 0.2$. Both computed in the full dynamic model. Net present values, aggregated across markets; market quantity in tonnes, all other panels in euros. The dashed vertical line marks the social cost of carbon (€150/tonne).

The reason is that output-based rebating can maintain high production and investment incentives. Figure 16b shows that, even at the higher tax rate, the production subsidy from rebating substantially reduces output suppression compared with a tax. At the optimal configuration, domestic production falls by only 2% relative to no carbon pricing, compared with 50% under the optimal tax alone (Table 4). In addition, this effective output subsidy increases the investment incentives even for a given tax rate, since the expected direct cost savings are higher due to the increased output (Figure 16c). Thereby, the OBA policy dominates the simple tax through two margins. Firstly, for a given tax, the investment incentive is higher. Second, it allows for a higher tax rate, bridging an additional part of the investment gap, whilst maintaining a higher production level through the production subsidy, even for low-carbon firms. The rebate also outperforms the subsidy and tax combination. Unlike the implicit production subsidy of a low-carbon tax, the rebate does not fade as firms approach emission neutrality, and even the cleanest firms receive a production subsidy.

However, at the estimated optimal output-based rebating configuration, the carbon revenues are negative (Figure 16d), indicating that this policy does require public funds. Relative to the optimal tax alone, the policy requires €20.8 billion of additional public funds, and its welfare gain over the individual tax remains up until a maximum cost of public funds of 20%.

6.4 Carbon Border Adjustment

Since 2026, the EU Carbon Border Adjustment Mechanism (CBAM) extends the carbon price to the emissions embedded in imported cement, while free allocation is gradually phased out. To evaluate this policy, I let importers pay the domestic carbon price on their embedded emissions, at the baseline emission intensity of $\eta^{Imports} = 0.63$. Thereby, a higher domestic carbon price no longer shifts production towards unregulated imports.

With a CBAM in place, the optimal carbon tax rises from €135/tonne to €175/tonne, and thus lies above the social cost of carbon (Table 4). This is because the CBAM removes the import channel. Raising the carbon price no longer shifts production and rents towards unregulated imports, but instead charges importers for their embedded emissions. Compared with the optimal tax alone, imports fall by 60%, and domestic producers regain market share. Despite the higher tax, consumer surplus therefore remains almost unchanged, while investment costs rise from €8.6 billion to €14.9 billion. Since the carbon tax remains the only instrument to address firms' insufficient investment incentives, the welfare-maximising tax exceeds the social cost of carbon. As a result, the CBAM reduces emissions, including those from imports, by a further 52.2%, and increases welfare by 25.7% relative to the optimal tax alone. However, around four-fifths of this gain is a transfer from importers.

Finally, I combine the CBAM with the complementary instruments. Without a border adjustment, imports cap the domestic price firms can charge, so a high carbon price erodes their margins and market share, and investments cannot be refinanced with sufficient additional profits. Since the CBAM subjects importers to the same carbon price, domestic producers can now better pass the carbon cost on to the market price. Consequently, profits remain high despite the higher tax, so that investments can be recouped, and the carbon price itself provides strong enough investment incentives. A subsidy therefore delivers only a small additional welfare gain of 0.4%, compared with 4.8% without a border adjustment. The output-based rebate still raises consumer surplus by counteracting the output contraction due to market power, but its additional welfare gain drops to 2.1%, compared with 8.5% without a border adjustment, and it slightly increases emissions by 5.9%. Thus, a substantial share of the welfare gains of the output-based rebate stems from protecting domestic producers against unregulated imports, a role the CBAM fulfils more directly. These results suggest that phasing out free allocation alongside the introduction of the CBAM forgoes only limited welfare gains.

Table 4: Outcomes under the Welfare-Maximising Configuration of Each Policy

	Without CBAM					With CBAM			
	Laissez-faire	Static optimal tax	Optimal tax	Tax + subsidy	Tax + rebate	Static optimal tax	Optimal tax	Tax + subsidy	Tax + rebate
<i>Policy</i>									
Carbon tax (€/t)	0	90	135	70	165	150	175	170	160
Adoption subsidy (%)	–	–	–	60	–	–	–	20	–
Rebate benchmark (t CO ₂ /t)	–	–	–	–	0.20	–	–	–	0.15
<i>Production (Mt)</i>									
Domestic production	673.9	399.1	335.6	604.7	661.1	394.6	408.0	513.3	651.4
Imports	107.8	119.8	123.9	110.5	109.3	66.9	49.8	37.5	29.2
Domestic emission intensity (t/t)	0.565	0.352	0.168	0.096	0.069	0.128	0.081	0.048	0.076
<i>Emissions (Mt CO₂)</i>									
Domestic emissions	380.9	140.5	56.4	57.9	45.6	50.3	32.8	24.9	49.6
Emissions from imports	67.9	75.5	78.1	69.6	68.9	42.1	31.4	23.6	18.4
Total emissions	448.8	215.9	134.4	127.6	114.4	92.5	64.3	48.5	68.0
<i>Domestic welfare (€bn)</i>									
Consumer surplus	79.23	64.45	60.49	75.64	78.07	60.44	60.00	65.78	73.40
Domestic profits	18.32	6.58	10.08	14.88	29.84	13.26	15.46	17.38	27.57
Investment cost	–0.25	3.50	8.56	23.61	23.87	12.46	14.89	21.70	23.34
Carbon revenue, domestic	0.00	12.64	7.61	4.06	–13.24	7.55	5.75	4.23	–6.89
Carbon revenue, CBAM	0.00	0.00	0.00	0.00	0.00	6.32	5.50	4.02	2.94
Damages, domestic emissions	57.14	21.07	8.46	8.69	6.84	7.55	4.93	3.73	7.45
Damages, emissions from imports	10.18	11.32	11.71	10.44	10.33	6.32	4.71	3.54	2.76
Domestic welfare	30.47	47.79	49.45	51.83	53.63	61.23	62.18	62.43	63.47
Net public funds	0.00	12.64	7.61	–10.11	–13.24	13.87	11.24	3.91	–3.94
Importers' profits	7.31	10.05	11.18	7.91	7.75	2.00	1.28	0.88	0.47

Notes: Social cost of carbon of €150/tonne. Discounted values, aggregated across markets. Domestic welfare equals consumer surplus, domestic profits and carbon revenue, minus investment cost and damages. Importers' profits are not part of domestic welfare. The competitive fringe does not receive a rebate.

7 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's 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. Accordingly, to meet its climate neutrality targets without triggering substantial deindustrialisation, EU climate policy aims to reduce emissions primarily by adopting green production technologies. Recognising that market concentration and competition with imports not subject to the same carbon price are key features 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. For a social cost of carbon of €150/tonne, I estimate an optimal carbon tax of €135/tonne, compared with €90/tonne implied by a static oligopoly analysis, as the option to invest in reducing the emission intensity of production eases the burden of carbon prices on domestic production. Still, market power, as well as the loss of market share and rents to unregulated imports, keeps the optimal carbon price under a carbon tax alone below the Pigouvian level, and firms invest substantially less than would maximise welfare, as they do not internalise the gains in consumer surplus from their investments. A green investment subsidy can deliver welfare gains of up to 4.8%, and an output-based rebate, which allows a production subsidy even for clean firms, of up to 8.5%. A carbon border adjustment changes this picture. Once importers pay the same carbon price, domestic producers can pass on a higher carbon price without losing market share, and earn sufficient margins to recoup their investments. The additional gain from a subsidy then largely vanishes, and that from an output-based rebate falls to 2.1%.

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. 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 designing such regulations. Combining a carbon tax with green investment subsidies or output-based rebates can reduce emissions further than a carbon tax alone, while limiting output losses in the domestic industry, but these gains come at a cost to public funds, which policy-makers need to weigh against the welfare gains. Moreover, a substantial part of these gains stems from protecting domestic producers against unregulated imports, a role a carbon border adjustment

fulfils more directly. These results suggest that phasing out free allocation alongside the introduction of the CBAM forgoes only limited welfare gains.

References

- Buchanan, J. (1969). External Diseconomies, Corrective Taxes, and Market Structure. *American Economic Review*, 59(1), 174–77.
- Hansen, L. P. (1982). Large Sample Properties of Generalized Method of Moments Estimators. *Econometrica*, 50(4), 1029–1054.
- 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>
- Newey, W. K., & McFadden, D. (1994). Large Sample Estimation and Hypothesis Testing. In R. F. Engle & D. L. McFadden (Eds.), *Handbook of Econometrics* (pp. 2111–2245, Vol. 4). Elsevier.
- 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](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](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](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](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>
- Mansur, E. T. (2007). Do oligopolists pollute less? evidence from a restructured electricity market. *The Journal of Industrial Economics*, 55(4), 661–689. <https://doi.org/10.1111/j.1467-6451.2007.00325.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].
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*.
- Leslie, G. (2018). Tax induced emissions? Estimating short-run emission impacts from carbon taxation under different market structures. *Journal of Public Economics*, 167, 220–239. <https://doi.org/10.1016/j.jpubeco.2018.09.010>
- Marie, M. J., Gaied, K. S., & Ajel, A. R. (2018). Wireless Network Control System for Electrostatic 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>
- Aguirregabiria, V., Collard-Wexler, A., & Ryan, S. P. (2021). Dynamic games in empirical industrial organization. In *Handbook of Industrial Organization* (pp. 225–343, Vol. 4). Elsevier. <https://doi.org/10.1016/bs.hesind.2021.11.004>
- 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. (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. *Ministère 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. (2023). Directive (EU) 2023/959 of the European Parliament and of the Council of 10 May 2023 amending Directive 2003/87/EC establishing a system for greenhouse gas emission allowance trading within the Union and Decision (EU) 2015/1814 concerning the establishment and operation of a market stability reserve for the Union greenhouse gas emission trading system [CELEX: 32023L0959].
- 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>
- Preonas, L. (2024). Market Power in Coal Shipping and Implications for U.S. Climate Policy. *Review of Economic Studies*, 91(4), 2508–2537. <https://doi.org/10.1093/restud/rdad090>
- 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 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>
- Eurostat. (2026). Greenhouse gas emission accounts by NACE rev. 2 activity [Accessed: September 28, 2026].

Appendix

A Data Sources and Descriptives

Table 5: Summary Statistics by Variable and Market

Variable	Market	Mean	SD	P25	P50	P75	Max
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
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 6: Overview of Variables and Data Sources

Data Point	Source	Dataset	Level	Access
Investment	INSEE	ANTIPOL	Plant	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	CASD
Revenue	INSEE	EAP	Firm	CASD
Prices	INSEE	EAP	Firm	CASD
Demand	France Ciment	L'Essentiel	Market	Public
Imports	France Ciment	L'Essentiel	Market	Public

Table 7: Plant Closures and Transformations

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 8: Capacity Expansion Events

Firm	Plant	Market	Date	Capacity Expansion
ECOCCEM	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 9: Entry of Fringe Firms

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
ECOCCEM	Dunkirk	NorthWest	Grinding	750	2018
Cem In EU	LesValences	SouthEast	Grinding	240	2021

Table 10: Investment Events Summary (sorted by Firm and Plant)

Firm	Plant	Market	Type	M_EUR	CO2ta	CO2rate	Announced	Expansion
ECOCER	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

Notes: M_EUR denotes investment spending in million euros. CO2ta denotes the annual CO₂ emissions saved, in thousand tonnes per year. CO2rate denotes the implied reduction in emission intensity, in percent. All figures are as reported in company announcements and press coverage. Archived links to the articles are available upon request. Blank cells indicate that the figure was not reported.




Figure 17: Import Share by Year and Market

B Estimation

B.1 Static Estimation




Figure 18: Moment Prediction - Operating Margin per Market




Figure 19: Model-Implied Operating Margin per Firm




Figure 20: Moment Fit - Market Level Production



Figure 21: Moment Fit - Firm Level Production




Figure 22: Moment Fit - Market Level Prices




Figure 23: Ratio - Inferred versus Observed Production

B.2 Dynamic Estimation

Robustness to Carbon Price Expectations

Table 11: Dynamic Cost Parameter Estimates by Long-Term Carbon Price Expectation

Expectation (€/t)	λ_1	λ_2	λ_3
50	0.667 (0.445)	0.095 (0.021)	1.347 (0.256)
75	0.492 (1.597)	0.123 (0.030)	1.681 (0.397)
100	0.122 (2.587)	0.156 (0.047)	2.019 (0.664)
125	0.067 (3.092)	0.180 (0.055)	2.221 (0.796)
150	0.067 (3.858)	0.201 (0.067)	2.365 (0.955)
175	0.067 (4.101)	0.225 (0.069)	2.544 (0.984)

Notes: Standard errors in parentheses are GMM standard errors under efficient weighting. For expectations of €125/t and above, the estimate of λ_1 lies close to its lower bound of zero.

GMM Moments and Standard Errors

Moment Conditions. I estimate the dynamic cost parameters $\lambda = (\lambda_1, \lambda_2, \lambda_3)$ using a nested fixed-point GMM estimator. For each parameter guess, I solve the single-agent dynamic problem of each firm, given the reduced-form competitor responses, and compute the implied conditional choice probabilities over investment levels. From these, I compute the deviation between model and data for each firm-period observation. Averaging these deviations across all observations gives the vector of mean moment deviations, $\bar{g}(\lambda)$. For the weighting matrix, $\hat{W}$, I use the inverse of the variance-covariance matrix of the mean data moments, which I estimate from 5,000 bootstrap draws of the investment panel.

Asymptotic Variance. To compute standard errors, I first compute the Jacobian of the mean moments at the estimated parameters, $D = \partial \bar{g}(\hat{\lambda}) / \partial \lambda'$, using central finite differences. Further, let $\hat{\Omega}$ denote the heteroskedasticity-robust covariance matrix of the moments, computed from the centred moment deviations of each observation. For a given weighting matrix W , the asymptotic variance of the GMM estimator is then given by:

$$\widehat{\text{Avar}}(\hat{\lambda}) = (D'WD)^{-1}(D'W\hat{\Omega}WD)(D'WD)^{-1}/N \quad (28)$$

where N denotes the number of firm-period observations. If the weighting matrix is efficient, i.e., $W = (\Omega/N)^{-1}$, this expression simplifies to $(D'WD)^{-1}$ (Hansen, 1982; Newey & McFadden, 1994).

Since I construct $\hat{W}$ as the inverse of the bootstrapped covariance matrix of the mean moments, which estimates Ω/N , I treat it as efficient and use the simplified expression for the standard errors reported in the main text. In the estimation, I search over a transformed, unconstrained version of the parameters. I then map the estimates back into the cost function parameters reported in Table 3 and compute their standard errors using the delta method.

Robustness: Sandwich Standard Errors. Treating $\hat{W}$ as efficient is an assumption. To check whether it matters for inference, Table 12 additionally reports standard errors based on the full sandwich formula, which do not require the weighting matrix to be efficient. Both sets of standard errors are similar in magnitude. λ_2 and λ_3 are significant at conventional levels in both cases, whereas λ_1 is imprecisely estimated under either approach.

Table 12: Dynamic Cost Estimates: Efficient-Weighting vs. Sandwich Standard Errors

	λ_1	λ_2	λ_3
Estimate	0.067	0.180	2.221
Std. error (efficient weighting)	3.092	0.055	0.796
95% CI (efficient weighting)	[-5.993, 6.127]	[0.071, 0.289]	[0.661, 3.780]
Std. error (sandwich)	3.577	0.063	0.993
95% CI (sandwich)	[-6.944, 7.078]	[0.056, 0.304]	[0.275, 4.167]

Notes: Efficient-weighting standard errors assume that the bootstrapped weighting matrix is efficient and are the ones reported in the main text. Sandwich standard errors are based on the full GMM variance formula and do not require this assumption.

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 $x = (x_1, \dots, x_N) \in \mathbb{R}^N$, where N denotes the number of dimensions of the state space (here, N is the number of firms per market, and x denotes the vector of emission intensities). For an off-grid state vector x , I define lower and upper grid points surrounding x for each firm as x_k^- and x_k^+ . Then, one can determine the position of any point within the state space relative to each firm's surrounding corner states. For dimension k , this reads as:

$$t_k = \frac{x_k - x_k^-}{x_k^+ - x_k^-} \in [0, 1] \quad (29)$$

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

$$\tilde{V}(x) = \sum_{c \in \{0,1\}^N} w_c(t) V(x^{(c)}) \quad (30)$$

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:

$$w_c(t) = \prod_{k=1}^N \begin{cases} 1 - t_k, & \text{if } c_k = 0 \\ t_k, & \text{if } c_k = 1 \end{cases} \quad (31)$$

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 24 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 $(0,0)$, $(10,0)$, $(0,10)$ and $(10,10)$. 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:

$$w_{00} = (1 - t_1)(1 - t_2) = (1 - 0.25)(1 - 0.5) = 0.75 \times 0.5 = 0.375,$$

$$w_{10} = t_1(1 - t_2) = 0.25 \times 0.5 = 0.125,$$

$$w_{01} = (1 - t_1)t_2 = 0.75 \times 0.5 = 0.375,$$

$$w_{11} = t_1 t_2 = 0.25 \times 0.5 = 0.125.$$

The interpolated value is then obtained as:

$$\tilde{V}(2.5, 5) = 0.375 V(x^{(0,0)}) + 0.125 V(x^{(1,0)}) + 0.375 V(x^{(0,1)}) + 0.125 V(x^{(1,1)}) \quad (32)$$





Figure 24: 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 damping 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 damping 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 the maximum iteration count is reached without convergence, the solver stops and accepts the last dampened iteration. For the paper's estimated parameters and specifications, this is rarely the case. It can also, for example, occur in some robustness checks with different cost parameters, most notably when the private shock variance used is insufficient.

D Counterfactual Outcomes


D.1 Carbon Tax




(a) Welfare




(b) Consumer surplus




(c) Profit



(d) Carbon revenue




(e) Damages




(f) Investment cost

Figure 25: Tax: Welfare Components


Notes: Full Model: dynamic model with abatement investment; Static Oligopoly: model without the investment margin. Net present values, aggregated across markets. The dashed vertical line marks the social cost of carbon (€150/tonne), where applicable.




(a) Quantity produced



(b) Market quantity



(c) Carbon leakage




(d) Investment cost (social optimum)


Figure 26: Tax: Quantities and Further Outcomes

Notes: Full Model: dynamic model with abatement investment; Static Oligopoly: model without the investment margin. Net present values, aggregated across markets. The dashed vertical line marks the social cost of carbon (€150/tonne), where applicable.


D.2 Technology Cost Subsidy




(a) Welfare




(b) Consumer surplus




(c) Profit



(d) Carbon revenue




(e) Damages




(f) Investment cost

Figure 27: Subsidy: Welfare Components


Notes: Carbon tax fixed at €70/tonne; the horizontal axis shows the subsidy as a share of investment cost. Net present values, aggregated across markets. The dashed vertical line marks the social cost of carbon (€150/tonne), where applicable.




(a) Quantity produced




(b) Market quantity



(c) Carbon leakage



(d) Subsidy cost




(e) Net public funds


Figure 28: Subsidy: Quantities and Further Outcomes

Notes: Carbon tax fixed at €70/tonne; the horizontal axis shows the subsidy as a share of investment cost. Net present values, aggregated across markets. The dashed vertical line marks the social cost of carbon (€150/tonne), where applicable.


D.3 Output-Based Rebate




(a) Welfare




(b) Consumer surplus




(c) Profit



(d) Carbon revenue




(e) Damages




(f) Investment cost

Figure 29: OBA: Welfare Components


Notes: Tax: carbon tax alone; OBA: carbon tax with an output-based rebate at a benchmark of $\eta^* = 0.2$. Net present values, aggregated across markets. The dashed vertical line marks the social cost of carbon (€150/tonne), where applicable.



(a) Quantity produced



(b) Market quantity




(c) Carbon leakage


Figure 30: OBA: Quantities and Further Outcomes

Notes: Tax: carbon tax alone; OBA: carbon tax with an output-based rebate at a benchmark of $\eta^* = 0.2$. Net present values, aggregated across markets. The dashed vertical line marks the social cost of carbon (€150/tonne), where applicable.


D.4 Carbon Border Adjustment




(a) Welfare




(b) Consumer surplus




(c) Profit



(d) Carbon revenue




(e) Damages




(f) Investment cost

Figure 31: CBAM: Welfare Components


Notes: Full Model: dynamic model with abatement investment; Static Oligopoly: model without the investment margin. Net present values, aggregated across markets. The dashed vertical line marks the social cost of carbon (€150/tonne), where applicable.



(a) Quantity produced



(b) Market quantity



(c) Emissions from imports

Figure 32: CBAM: Quantities and Further Outcomes

Notes: Full Model: dynamic model with abatement investment; Static Oligopoly: model without the investment margin. Net present values, aggregated across markets. The dashed vertical line marks the social cost of carbon (€150/tonne), where applicable.

D.5 Sensitivity to SCC

Table 13: Sensitivity to the Social Cost of Carbon

Social cost of carbon (€/t)	75	100	125	150
<i>Without CBAM</i>				
Optimal tax (€/t)	45	65	110	135
Static tax (€/t)	25	50	80	90
welfare vs. optimal tax	−0.7%	−0.6%	−1.2%	−3.4%
Tax + subsidy (€/t, %)	15, 80%	20, 80%	40, 70%	70, 60%
welfare vs. optimal tax	+3.4%	+4.5%	+4.6%	+4.8%
Tax + rebate (€/t, t/t)	80, 0.30	105, 0.30	135, 0.25	165, 0.20
welfare vs. optimal tax	+3.5%	+5.6%	+6.6%	+8.5%
<i>With CBAM</i>				
Optimal tax (€/t)	85	115	170	175
welfare vs. optimal tax without CBAM	+6.6%	+12.3%	+19.4%	+25.7%
Tax + subsidy (€/t, %)	60, 35%	115, 5%	155, 10%	170, 20%
welfare vs. optimal tax with CBAM	+0.7%	+0.1%	+0.2%	+0.4%
Tax + rebate (€/t, t/t)	95, 0.30 [†]	115, 0.25	140, 0.20	160, 0.15
welfare vs. optimal tax with CBAM	+5.1%	+4.4%	+2.6%	+2.1%

Notes: Welfare-maximising configurations under domestic welfare. [†]Optimum at the boundary of the search grid.