When to go Green? Firm dynamics & Clean Technology Adoption *

Bas Gorrens

KU Leuven

Job Market Paper
(2025-2026)

January 14, 2026

Abstract

Abstract: Carbon pricing is a central policy instrument for reducing emissions, but governments face a trade-off: faster decarbonization can raise output losses and carbon leakage, while gradual implementation slows emission reductions. This paper studies how EU carbon policies have shaped firms' adoption of abatement technologies and identifies the optimal trajectory to reach the EU's 2050 net zero target, particularly in a unilateral context. I develop a dynamic heterogeneous-firm model in which forward-looking manufacturing firms choose when to adopt discrete abatement technologies under a gradually tightening carbon price. I estimate it using panel data on EU ETS firms from 2005-2019. The model rationalizes the low carbon prices of the 2010s as a consequence of gradual policy and firm anticipation. Emission reductions arise mainly from large, productive, and initially polluting firms. Anticipation of future tightening mitigates half of the short-run output losses in 2025 and two-thirds by 2050, keeping overall output losses below 2%. A moderately faster tightening could cut cumulative emissions by 15% at an additional cost of only 0.11% of output. Finally, because firms anticipate future policy changes, unilateral and global carbon pricing yield nearly identical effects on domestic output and carbon leakage.

Keywords: trade and environment, technology adoption, firm decisions, climate policy, carbon leakage.

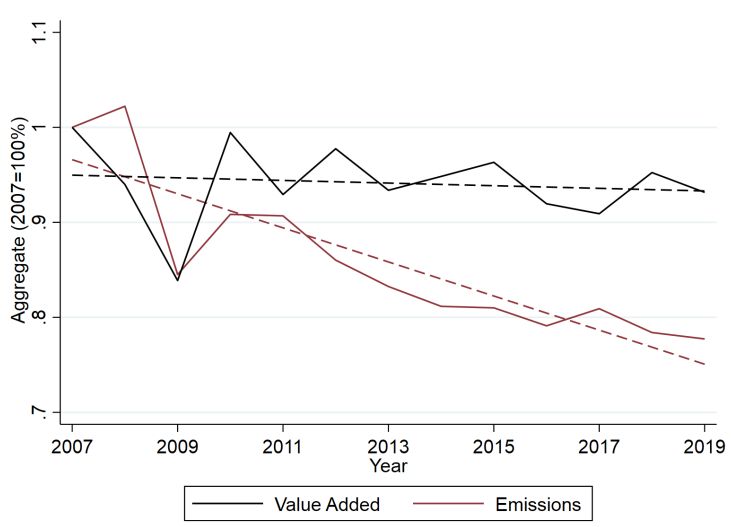
JEL Codes: F18, F64, L1, Q56.

*I am deeply grateful to Johannes Boehm, Clement de Chaisemartin, Sharat Ganapati, Ahmad Lashkaripour, Farid Farrokhi, Estelle Cantillon, Sareh Vosooghi, Jo Van Biesebroeck, Mirabelle Muuls, Mar Reguant, Florian Forsthuber, Charlotte De Canniere, and Joep Konings for their valuable discussions, insights, and comments. This paper has also greatly benefited from feedback received at the EEA-ESEM Conference (Rotterdam), the 30th Dubrovnik Young Economist Seminar, the Workshop on European Integration, Economic Convergence and Productivity, the SASCA PhD Conference, the Royal Economic Society Conference (Birmingham), EAERE (Bergen), EARIE (Valencia), the Barcelona School of Economics, and various internal workshops and seminars.
Mail author: bas.gorrens@kuleuven.be.

1. Introduction

Achieving large emission reductions without harming economic performance is a central challenge for climate policy. This paper studies the mechanisms behind emission reductions and how governments can design carbon policies that foster clean technology adoption and minimize output losses - even when implemented unilaterally. Governments worldwide have introduced carbon taxes and cap-and-trade systems. Yet, despite their widespread adoption, these instruments remain controversial (Drupp et al., 2024). Environmental groups argue that carbon prices are too low to spur substantial investment in clean technologies, delaying the transition to net zero. Industry stakeholders, in contrast, argue that if policies tighten too quickly, firms lack time to adopt new technologies, eroding competitiveness, reducing domestic welfare, and potentially shifting emissions abroad.

Figure 1: Aggregate evolution manufacturing firms in EU ETS



Note: The figure shows the evolution of total emissions and value added for all manufacturing firms covered by the EU ETS since 2008. Values are normalized to 100% in 2007. The black line represents the deflated sum of total value added, the red line indicates the sum of annual emissions emitted by all firms in the baseline dataset. A full explanation of the dataset can be found in Appendix A. Firms whose number of regulated plants increased during this period are excluded. The dashed lines show a linear trend for each.

The European Union Emissions Trading Scheme (EU ETS) offers a particularly instructive case to demonstrate the implications of a gradual approach to achieve full decarbonization while limiting output losses. Established in 2005, it is the most comprehensive and stringent carbon market to date, combining an explicit commitment to net-zero emissions by 2050 with a gradually tightening cap on CO2 emissions designed to achieve a smooth decoupling of emissions from output (Figure 1). The scheme is designed to leverage the forward-looking behavior of firms: credible long-run targets encourage early adoption of cleaner technologies, while gradual tightening allows firms time to adjust and maintain competitiveness. Determining the

optimal pace of tightening depends on the distribution of polluting firms, their adoption costs, and the degree of international spillovers (≈ carbon leakage) that influence the domestic burden of policy.

This paper develops a dynamic discrete choice model in which heterogeneous, forward-looking firms decide when to adopt novel clean technologies in response to an ever-tightening carbon policy. They operate in a quality-ladder environment for clean technologies. The novelty of the model lies in integrating two key firm characteristics into the setting of an ever-tightening carbon policy. First, firms are forward-looking: they anticipate future policy tightening and adjust their adoption decisions accordingly. Second, firms differ in productivity and existing abatement capacity, causing some to adopt cleaner technologies earlier than others. This heterogeneity generates a catch-up mechanism in which the most productive and most polluting firms adopt first. The interaction of these two key characteristics explains how gradual policy tightening can achieve substantial emission reductions while keeping output losses modest.

Using the insights of the model, the paper assesses how firms respond to the EU ETS by examining panel data covering all manufacturing firms in the system from 2005 to 2019. The results show that emission reductions are primarily driven by more productive and more polluting firms, while reallocation through market exit, entry, or shifts in market shares plays only a minor role.

To quantify aggregate effects and assess potential distortions from carbon leakage, I embed the firm-level model in a two-country trade framework representing the EU and the rest of the world (ROW). The model features general equilibrium interactions in both the goods and (domestic) emissions markets: the goods market captures substitution toward cleaner domestic products and foreign imports, while the emissions market determines the carbon price through market clearing. Firms are price takers on the carbon market. The economies are initially symmetric before the EU introduces a gradually tightening emissions cap aimed at reaching net zero by 2050. This policy path temporarily shifts the economy from its balanced growth trajectory, but as firms adopt cleaner technologies, it converges toward a new equilibrium in which all firms operate at the clean technology frontier.

The state variables of firms in the model - productivity and abatement technology - are identified using Phase I ETS data (2005 - 2007), a period with limited policy stringency that provides a clean baseline (Cantillon and Slechten, 2023; Colmer et al., 2020). Adoption decisions are inferred from structural changes in firms' abatement technology, which are empirically associated with increased capital investment rather than reductions in output, indicating that adoption, rather than contraction, drives emission reductions. Consistent with policy incentives, clean-technology adoption intensifies when carbon allowance prices are high. The joint distribution of productivity and abatement technology, needed to capture general equilibrium effects, is recovered non-parametrically. The model is solved by backward induction, and the remaining parameters governing adoption costs are estimated via maximum likelihood.

Before analyzing aggregate outcomes, the model is first validated at the firm level to demonstrate that aggregate dynamics are driven mainly by highly productive, polluting incumbents that adopt cleaner technologies more frequently. The model closely replicates these observed patterns in both adoption decisions

and outcomes. It highlights the importance of the adoption channel as the underlying mechanism by which firms become greener rather than reallocation between firms.

The model endogenously generates a carbon price path that mirrors the gradual increase in spot prices observed in the EU ETS. The equilibrium price emerges from market clearing in the emissions market, where firms' demand for allowances interacts with a tightening emissions cap. This demand is shaped by the evolution in their adoption costs and conditions in the output market.

Under the current EU ETS trajectory, domestic emissions in the European Union are set to decline to about 10% of their 2005 level by 2050. The model projects that following this path entails moderate economic costs, with domestic demand (sales) falling by roughly 0.5% (1%) in 2025 and 1.9% (4.5%) in 2050, before the economy gradually converges back to its balanced growth path. These effects operate through two channels: a direct output channel, where stricter policy temporarily curbs production, and an adoption channel, where the adoption of cleaner technologies decouples output from emissions and mitigates losses. In the early stages, the reduction in output stems almost entirely from the direct output channel. Over time, the mitigating adoption channel becomes increasingly important, contributing around 50% of reductions in 2025 and 65% in 2050, and ultimately accounting for 100% once the economy reaches net zero. By anticipating future policy tightening, firms shift toward cleaner technologies early, which flattens the aggregate implications of the carbon policy on domestic prices. This preemptive adoption limits costly reallocation between firms and reduces the scope for carbon leakage, making the overall macroeconomic impact of the policy relatively modest.

The proposed framework enables analysis of dynamic, firm-level responses to carbon policies that extend beyond the static setting commonly used in the trade literature. I consider two counterfactual instruments and their combination: (i) a carbon border adjustment mechanism (CBAM) that extends coverage to emissions-intensive imports, and (ii) revenue recycling through adoption subsidies instead of lump-sum transfers. The CBAM reduces foreign cost advantages, which expands domestic firms' market share. This higher market share raises carbon demand, pushing up carbon prices and accelerating adoption, though overall efficiency gains remain modest as the adoption decisions happen based on firm-level expectations. In contrast, adoption subsidies directly shift expected returns to investment in clean technologies, accelerating diffusion and easing the long-run burden of carbon policy.

The main finding of this paper arises from applying the quantitative model to determine the optimal carbon price trajectory that achieves net zero by 2050 while minimizing output losses. Using the estimated structure and a social cost of carbon of $\lambda = 139€/\text{tonne}$, I compare the optimal unilateral policy to a coordinated global scenario. In the optimal unilateral case, cumulative domestic emissions fall by 15% by 2050, with only a 0.11% additional output loss and a 1% increase in adoption. The global optimum yields a 12% reduction for a 0.18% output loss, while adoption responses increase with 0.8%.

In the optimal scenario, the government places slightly greater weight on emissions relative to output, resulting in a faster tightening at the start. The effectiveness of the unilateral policy stems from its initially

higher yet more predictable carbon price trajectory, which anchors firms' expectations and triggers earlier adoption of clean technologies. Because firms are forward-looking, they invest in anticipation rather than reacting to shocks, which limits output losses. Consequently, the global policy outcome differs little from the unilateral one: the EU's large share of the world economy and the gradual tightening of its emissions cap dampen reallocation and minimize carbon leakage. In the global scenario, the substitution channel effectively shuts down, as demand can no longer shift toward cheaper foreign goods, while the marginal benefit of abatement rises when all countries participate. As a result, the optimal global trajectory is slightly slower at first than the unilateral path. Finally, I show the trajectory for different values of λ commonly used in the literature.

Finally, instead of finding the optimal trajectory given a certain cost of carbon λ, I invert the analysis and argue that the government has taken into account a certain cost of carbon when designing this optimal trajectory of carbon. Put differently, the model can back out the social cost of carbon internalized by the EU in selecting the current emissions trajectory over all alternative paths. The social cost of carbon that justifies this trajectory is $\lambda = 81€/\text{tonne}$.

1.1 Related literature

This paper relates to four main strands of the existing literature: (i) trade and environment with heterogeneous firms, (ii) dynamic models of firm behavior, (iii) empirical analyses of carbon policies, (iv) and macroeconomic studies of long-run climate policy.

The first strand concerns the trade-and-environment literature on abatement in models with heterogeneous firms (Copeland and Taylor, 2004; Shapiro and Walker, 2018; Forslid et al., 2018; Cherniwchan and Najjar, 2022, ...), which build on the seminal framework of Melitz (2003). These studies emphasize that large firms account for a disproportionate share of aggregate emissions yet tend to be relatively less polluting. My paper builds directly on this framework but departs in one crucial respect: it models abatement technology adoption as an irreversible, forward-looking decision. This dynamic formulation allows firms to anticipate tightening carbon constraints and adjust their technology choice over time, rather than reacting myopically to current prices.

Building on this literature, I next connect to research that highlights the role of dynamic frictions in shaping firm and trade outcomes (Atkeson and Burstein, 2010; Alessandria et al., 2021a,b). By modeling the dynamic decision explicitly, this paper extends that body of work to a setting in which emission outcomes evolve endogenously along the transition to a low-carbon economy. The novelty lies in two aspects: first, the explicit modeling of gradual policy implementation, and second, the incorporation of an environmental block into the analysis.

Moreover,on the trade margin, this analysis connects to a large body of papers on carbon leakage and the design of optimal carbon policy in a unilateral setting (Kortum and Weisbach, 2021; Farrokhi and Lashkaripour, 2022; Fowlie and Reguant, 2022; Fontagne and Schubert, 2023; Shapiro, 2021; Bohringer et al.,

2017); for a comprehensive overview, see Felbermayr et al. (2025). This paper introduces a new dimension to optimal policy design: the dynamic path toward decarbonization. This path shapes firms' forward-looking technology choices, smooths aggregate price adjustments, and ultimately mitigates carbon leakage.

This paper is most closely related to Fowlie et al. (2016) and Ryan (2012), who study technology adoption and market dynamics for forward-looking firms in the U.S. cement industry under alternative policy regimes. My framework departs in three key dimensions. First, both the output and emissions markets clear, allowing for reallocation between firms and a comprehensive welfare analysis. Second, I introduce a two-country trade structure, making carbon leakage an endogenous outcome. Finally, I incorporate an emissions cap that tightens over time toward net zero, generating an endogenous carbon price path and enabling evaluation of dynamic policy trajectories, including the optimal unilateral path.

While the preceding work is largely structural, my paper also extends and complements the existing empirical literature that analyzes the ex-post impacts of carbon policies, particularly the EU ETS (Dechezlepretre and Sato, 2017; Martin et al., 2016; Verde, 2020; Colmer et al., 2020; De Jonghe et al., 2020; Borghesi et al., 2020; Dechezlepretre et al., 2018). Prior work has shown that higher abatement costs incentivize firms to reduce emissions and underscores the disproportionate contribution of large emitters to the overall decline in manufacturing sector emissions. This study provides a detailed examination of the underlying mechanisms and disentangles the contribution of abatement intensity from changes in total output from which counterfactual inference is possible.At the aggregate level, this study connects to work decomposing the evolution of clean technology (Najjar and Cherniwchan, 2021; Rottner and von Graevenitz, 2024; Shapiro and Walker, 2018; Barrows and Ollivier, 2018; Brunel and Levinson, 2021), which generally finds that within-firm emission reductions dominate effects from firm entry, exit, or market-share shifts. My paper emphasizes the forward-looking nature of the firm, together in combination with firm heterogeneity as a driving factor and quantifies these empirical decompositions using a structural model, which allows me to evaluate welfare implications.

Finally, by defining a dynamic carbon policy, this study draws inspiration from the literature on long-term macroeconomic implications of climate change and carbon policy. Building on the foundational work of Nordhaus (1992) and subsequent contributions (Golosov et al., 2014; Acemoglu et al., 2012; Hassler et al., 2021; Hemous, 2016, ...), it examines how carbon policy can guide the economy toward a decarbonized state (see Bilal and Stock (2025) for an overview). In particular, it relates to studies such as Campiglio et al. (2022) and van der Ploeg and Rezai (2020), which emphasize capital adjustment costs (as in Cooper and Haltiwanger, 2006) as a key channel through which climate policies affect output and welfare. I extend these aggregate frameworks by introducing a richer microstructure that captures firm heterogeneity, their forward-looking decisions, trade exposure, and dynamic technology adoption. This integration yields more granular predictions and a closer match to observed firm and industry data.

This paper is structured as follows. Section 2 examines how carbon policy shapes the adoption decision for firms. Section 3 describes the institutional setting of the EU ETS, and presents the associated dataset used

in this analysis. Section 4 characterizes the structural model of firm dynamics and estimates the dynamic discrete choice functions for the firms that participate in the EU ETS. Section 5 shows the aggregate results for the EU ETS and examines the implications of carbon leakage. Section 6 uses the estimated results from the earlier section and aims to compute the optimal carbon policy. Finally, Section 7 concludes.

2. Stylized model

To illustrate the conceptual trade-offs that forward-looking firms have to make when faced with an emission policy that targets their marginal costs, I introduce a simplified partial equilibrium dynamic model of heterogeneous firms that adopt increased levels of abatement technologies or decide to relocate abroad. This framework also presents notation that will be used throughout the paper. The objective of the government is to encourage firms to increase their adoption of clean technologies while minimizing the output loss incurred.

2.1 Climate policy

Government The government aims to encourage firms to adopt higher levels of clean technology. It can increase the price of carbon emissions either directly by taxing firm-level emissions (i.e. a carbon taxation) or indirectly by setting a limit on the supply of emissions (i.e. a Cap-and-Trade system), resulting in a price due to the laws of supply and demand.[1] Within the latter context, I assume firms are price-takers and take the market price as given. The annual price of 1 tonne of carbon can be written as a sequence $\zeta = \{\zeta_t, \zeta_{t+1}, ..., \zeta_T\}$.

Firms The firms are heterogeneous in their level of marginal cost $z_t \in \mathbb{R}_0^+$ and abatement technology $a_t \in \mathbb{R}_0^+$. The latter encompasses both the input composition and the capacity of the firm to emit less. By doing so, I adhere to one of the key aspects of the EU ETS, whereby firms are entrusted with the autonomy to determine the most effective measures to reduce their emission output. Emissions are a function of both underlying state variables, $e(z,a)$, and are weakly increasing in productivity and decreasing in clean technology.

Firms are presented with a discrete option in response to a governmental policy instrument: comply without making any operational changes or incur a sunk cost to adopt improved carbon-reducing technology, thereby mitigating the impact of the policy on their marginal costs. The value of a firm is represented by the Bellman equation:

[1] The Cap-and-Trade system differs from a carbon taxation in several ways. The latter imposes a tax on emissions directly, providing cost predictability but at a cost of flexibility (Metcalf, 2023; Goulder and Schein, 2013) The former leaves it to the market clearing of the carbon market to set a price. In the model used throughout the text, these two systems coincide as the model abstracts from any uncertainty as well as sectoral differences.

$$V_t(a_t, z_t|\zeta_t) = \max_{d_t \in \mathscr{D}} \{\pi(a_{t+1}, z_t, d_t|\zeta_t) + \beta \mathbb{E}[V_{t+1}(a_{t+1}, z_{t+1}|\zeta_{t+1}, a_t, z_t, d_t)]\}, \tag{1}$$

with static profits $\pi(a_{t+1}, z_t, d_t|\zeta_t)$ and discount factor $\beta \in (0,1)$. $\mathscr{D}$ denotes the set of $N+1$ possible actions with $\mathscr{D} = \{0, \chi_1, \chi_2, \ldots, \chi_N\}$. The decision to adopt is made at the beginning of the period. Profits are characterized by $\partial\pi/\partial a_t > 0$, and $\partial^2\pi/\partial^2 a_t < 0$ conditional on $\zeta > 0$ and $\partial\pi/\partial z_t > 0$.

In the business-as-usual case, there is no carbon policy ($\zeta_t = 0$). In this case, profits are unaffected: $\pi(a_t|\zeta = 0) = \pi$ for all a_t. Hence, they do not internalize other indirect effects of climate emissions in their adoption (as a function of their clean technology) - i.e. negative externalities, either through lower consumer value or increased aggregate risks may still be there. Firms will only adopt based on ζ.

Deciding to increase abatement levels ($d_t = \chi$) yields a cost $k(\chi)$ at the gain of $a_{t+1} = a_t\chi$ with a step-size increase $\chi > 1$.[2] The profit values are normalized to the case where the firm chooses to do nothing ($d_t = 0$) and can be stated as follows:

$$\pi(a_{t+1}, z_t, d_t|\zeta_t) = \begin{cases} \pi(a_t, z_t, 0|\zeta_t) & \text{if } d_t = 0 \\ \pi(a_t\chi, z_t, 0|\zeta_t) - k_t(\chi) & \text{if } d_t = \chi \end{cases} \tag{2}$$

for any of the N different choices of the step size parameter $\chi \in \{\chi_1, \ldots, \chi_N\}$. The parameter χ is chosen a priori. Since it is considered an investment, there is no uncertainty surrounding its implementation (i.e. $\mathbb{E}[\chi] = \chi$).

The DRS assumption of profit w.r.t. a_t together with the monotonicity assumption of the value V with π, generates a cutoff value of $\bar{a}_t$ for which $k_t(\chi) = V(\bar{a}_t\chi, z_t, 0|\zeta_{t+1}) - V(\bar{a}_t, z_t, 0|\zeta_{t+1})$. The cost of investment is time-dependent; intuitively, an improved installation with the same clean technology will become cheaper as time goes on.

For a firm with a given level of clean technology a_t, the effect of an increase in ζ_t can be summarized as follows: $\partial\pi_t/\partial\zeta_t < 0$ and $\partial^2\pi_t/\partial^2\zeta_t > 0$ as shown in Figure 2. In the figure, productivity is held constant.[3] Adjusting the value of ζ_t would raise the threshold at which firms become indifferent between adopting clean technology or not (from a^L to a^H).

The framework predicts convergence in emissions as highly polluting firms achieve more net gains from adopting novel levels of clean technology.

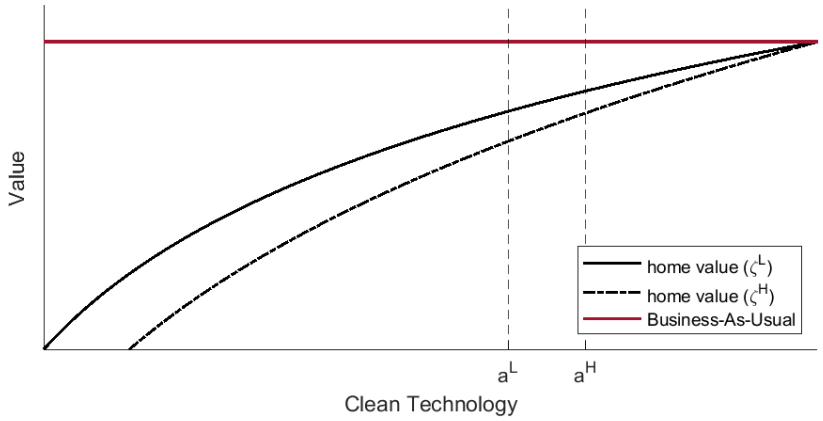
2.2 Optimal dynamic policy

Finally, this mechanism leads to the design of an optimal dynamic policy aimed at meeting a future emissions target. A high carbon price at the beginning encourages more firms to adopt new technologies, but

[2] Hence the model can be interpret as a quality ladder model on the clean technology margin.

[3] Heterogeneity in marginal costs means more productive firms can better absorb investment costs, facilitating abatement adoption. While abatement levels may converge, firm size measures like value added may diverge, as larger firms adopt clean technologies more easily.

Figure 2: Increase cost of Government policy ζ



Note: Graphical depiction of the implication of a carbon policy set by the home government for a firm with certain marginal cost z. The red line shows the present value of the business-as-usual case with $\zeta = 0$. The black line gives the value of the firm in the presence of a carbon taxation $V(a_t, z_t|\zeta_t)$. The dashed black line represents the value after an increase in the stringency of the policy $\zeta^L \to \zeta^H$. The thresholds line a^L shows the moment at which firms are indifferent between acquiring new abatement technologies or not $k_t(\chi) = V(a_t\chi, z_t, 0|\zeta_t^L) - V(a_t, z_t, 0|\zeta_t^L)$ for a constant level χ. An increase in taxation $\zeta^L \to \zeta^H$ shifts the value of the firm from the full line to the dashed line and subsequently pushes the decision thresholds to the right: $a^L \to a^H$.

it also reduces output significantly. In contrast, a lower initial carbon price results in fewer firms adopting these technologies, while causing less harm to demand. The optimal unilateral emissions policy therefore begins with a relatively low carbon price (or lenient cap) that gradually increases. This approach supports technology adoption over time without putting excessive pressure on output. This gets aggravated when the government proceeds unilaterally and firms face extensive competition from foreign firms exempt from carbon policy (i.e. carbon leakage).

As long as the government values maintaining output, a gradual increase in carbon pricing allows firms to continue producing while adapting to new technologies.[4] Appendix C illustrates this with a stylized two-period framework featuring three firm types and two alternative policy pathways. The model shows that the preferred route to decarbonization depends on (i) the distribution of firms and (ii) weight the government assigns to carbon. Intuitively, almost no highly polluting firms and a high social cost of carbon lets the government opt for the faster transition. Extending the model to include a foreign actor reveals that carbon leakage creates incentives for a more gradual policy: an aggressive domestic path reduces output competitiveness and shifts emissions abroad, thereby amplifying leakage.

To solve the model for a more realistic scenario and to conduct inference, a numerical solution is needed.

[4] As a corner solution, if the government values only output when achieving the target, it is sufficient to rely on firms' forward-looking behavior and guide them toward this constraint.

3. Institutional setting and reduced form evidence

To estimate the model, the paper leverages data from the European Union Emission Transaction system (ETS) - the largest carbon policy to date. This section provides a short summary on the EU ETS system - how it aims to reduce carbon leakage and how it is increasing in stringency over time. It also outlines the data sources employed in the model estimation and provides some reduced-form evidence for this *catch-up mechanism* where large polluting firms have reduced their emissions more, in the data.

3.1 The EU Emission Transaction System (ETS)

To achieve its objective of climate neutrality by 2050, the European Union has established a comprehensive regulatory framework to reduce greenhouse gas emissions. Central to this framework is the EU Emissions Trading System (EU ETS), the world's largest cap-and-trade scheme. It covers approximately 40% of all emissions in the EU. The following paragraphs provide a short overview of the trading system.[5]

The EU ETS sets a limit (cap) on the total (industry) emissions permitted each year.[6] The cap is designed to decrease emissions to approximately 10% of their 2005 level by 2050. By limiting the aggregate supply of emission allowances and enabling trading among firms, the system generates a market price for carbon and incentivizes cost-efficient abatement. This way, the governmental body leaves it to the individual firm to find the most effective way to become less polluting.

Since its inception in 2005, the EU ETS has evolved through four distinct phases: Phase 1 (2005-2007) marked its inception, but it faced challenges due to over-allocation of allowances, resulting in incomplete markets with low carbon prices. It has often been described as a phase of learning, both for the governmental agency and the participating firms (Cantillon and Slechten, 2023). Later, phase 2 (2008-2012) was the first true carbon reduction phase. It saw the expansion of the market and reductions in free allowances, and it was accompanied by an increase in the price of carbon emissions and a fully grown secondary emission market where agents could trade without the participation of the government. It has experienced a decrease of approximately 6.5% in aggregate emissions relative to phase 1. The cap on the total emissions was tightened in phase 3 (2013-2020) with a yearly linear reduction factor of 1.74% , and the supply of free allocated emission allowances reduced. This supply shifted from an ex-ante benchmark (grandfathering) to an industry-level benchmark (Martin et al., 2014). The final phase prior to reaching net zero, Phase 4 (2021-2030) is even more ambitious. It reduces the cap with a linear rate 2.2% of the 2005 emission level.[7]

[5] An extensive summary on the EU ETS can be found on the website of the European Union on Climate Action: EU Emissions Trading System (EU ETS)

[6] This cap applies only to direct (Scope 1) emissions, that is, greenhouse gases released from sources owned or controlled by the installation. Indirect emissions, those associated with purchased electricity (Scope 2) or other value-chain activities (Scope 3), are excluded from the policy.

[7] In 2023, a reform was applied that revised this trajectory. It would increase the linear reduction factor to 4.3% from 2023-27 and 4.4% from 2028 onward, essentially reducing the time from which the Emission cap trajectory would reach 0 from 2058 to 2038. This reform has changed the market stability reserve as well (Pahle et al., 2025). Since the reform was announced and implemented after the time frame considered in this analysis, firms had not yet incorporated it into their expectations. The policy change is therefore excluded from the baseline scenario in this study.

3.2 Data gathering and sources

The analysis makes use of the European Union Transaction Log (EUTL) as its primary data source, encompassing annual emissions data for each installation owned by firms participating in the ETS from 2005 to 2022. This dataset is then combined with firm-level financial information from the ORBIS database, maintained by Bureau Van Dijk. In doing so, I account for the considerations posed by Kalemli-Ozcan et al. (2015). The integration of these two datasets follows the matching procedure detailed in Letout (2022), and I manually correct. A more elaborate explanation on the data gathering process and the sources used can be found in Appendix A. The analysis is limited to the manufacturing sector only to ensure comparability between firms. As firms outside the manufacturing sector are structurally different, comparing them could lead to biased results.[8] Prices are deflated at the sector-country level using EUROSTAT data.

Table Table 1 reports summary statistics for manufacturing firms participating in the EU ETS. The dataset includes just over 3,000 firms; a relatively small number, however, the firms are notably large in scale. On average, firms employ 915 workers, generate an annual value added of 105€million, and emit 181 tonnes of CO2. The sector also appears highly stable, with annual entry and exit rates of less than 1%.

Table 1: Summary statistics

Variable	Mean	p10	p90
Number of employees	915	23	1490
Emissions (thousand tonne)	181	3.2	270.8
Value added (million Euro)	105.72	0.67	337.18
Number of installations	1.65	1	3
Extensive margin	Number of Firms	Entry rate	Exit rate
	3106	0.75%	0.65%

Note: Summary statistics for manufacturing firms included in the EU ETS. The sample is unbalanced and covers the period 2005-2019. The table reports the mean as well as the 10th and 90th percentiles for employment, emissions (in thousand tonnes), value added (in million euros) and the number of installations. All monetary values are deflated to 2015 levels. Since the ETS policy has become more stringent over time, the sample includes only firms whose ETS-covered installations have been present since the start of Phase III. This approach allows for relocation of installations between firms but does not increase the number of covered plants. Each firm is identified by a unique BVD-ID number. Entry (exit) refers to the first (last) financial observation of a firm in the dataset. Entry into the ETS itself is not considered unless it coincides with the first financial observation of the firm.

Throughout this analysis, several additional data sources are touched upon, such as the expected price of carbon certificates, oil and gas prices, and aggregate trade statistics. Data concerning the price of carbon

[8]For instance, the full list of firms that abide to the carbon policies contains electric and waste management companies among others, these are firms that do not trade abroad and for which the risk of carbon leakage is limited. They are considered different for the government as well and have received less carbon certificates. The aforementioned mechanism is therefore not necessarily transferable for these firms.

certificates are obtained from THOMPSON REUTERS. Trade data from WIOD tables are used for aggregate statistics.

3.3 Reduced form evidence

The stylized framework in Section 2 puts forward the hypothesis that large firms are becoming cleaner over time. To assess whether the observed evolution of the data aligns with this mechanism, I employ a local projections approach to track the emissions trajectory of ex-ante polluting incumbents. The analysis reveals that the most polluting firms reduce their emissions more significantly on average, while changes in their value added align more closely with Gibrat's law. Rather than a shift-share dynamic where greener firms have become larger, large incumbents have driven most of the results.

Next, one could see green entrants as representing an extreme case of potential shift-share evolution. I find that they tend to be smaller and less impactful overall, yet display higher emission intensity. Notably, their emission intensity does not converge to lower levels over time relative to incumbents, undermining the notion of a green-driven compositional effect..

To disentangle the relative contributions of within-firm improvements and market-share reallocation, I perform a decomposition following Melitz and Polanec (2015) and Barrows and Ollivier (2018) to the level of clean technology of the firm (weighted by value added as market share). The results show that within-firm improvements among continuing firms are the main driver of aggregate clean-technology gains. The reallocation term is modestly positive, indicating some market-share shift toward cleaner incumbents, while entry and exit effects remain small. These patterns confirm the earlier regression-based results and suggest that aggregate gains stem primarily from individual firm improvements rather than reallocation between firms. These decomposition results are in line with the decomposition of local pollutant emissions (Shapiro, 2021; Rottner and von Graevenitz, 2024; Najjar and Cherniwchan, 2021; Barrows and Ollivier, 2018).

Finally, to address the possibility of reallocation effects beyond the covered ETS firms in the EU, towards the non-covered manufacturing sector, I compare aggregate value-added growth between the ETS-covered firms and broader manufacturing firms using EUROSTAT data. The results indicate that ETS firms have experienced faster growth, suggesting that large-scale relocation into the non-ETS sector is unlikely to be driving the observed dynamics.

All the results can be found in more detail in Appendix B. Note that these findings are not conclusive but rather serve as suggestive evidence to guide the development of a structural framework.

4. Structural Estimation

The reduced-form analysis shows that the primary effects are driven by polluting incumbents. This raises the first question of the paper: to what extent can the observed decline in aggregate emissions be attributed

to the adoption of novel clean technology given the tightening carbon policy. To answer this, I revisit the discrete choice model from Section 2 and embed it within a general equilibrium asymmetric two-country trade model in which firms are heterogeneous in both their level of clean technology and productivity. I first lay out the building blocks of the model, and then turn to the estimation of the parameters, both outside the model and within its framework.

The model captures the transition dynamics from an initial state where firms operate in a setting in which the external costs of carbon emissions are not structurally internalized (i.e., no Pigouvian tax) to a future state where net-zero emissions have been achieved (Golosov et al., 2014) unilaterally by the home country. It is used to estimate the trajectory of the economy along this path toward a zero-carbon future and to infer counterfactual analyses that explore how different policy interventions could accelerate the transition.

4.1 Model

The firms' adoption decisions are modeled within a firm dynamics model with heterogeneous firms that trade, that extends the seminal contributions of Melitz (2003) and Atkeson and Burstein (2010) to a setting where climate policy is taken into account (Copeland and Taylor, 2004; Shapiro and Walker, 2018). Firms operate in two countries, the home country h, and the foreign country f, which are ex-ante symmetric. They provide their domestic market d and export to the foreign market x. Asymmetry arises solely through unilateral environmental policy differences implemented by the home country. The economy includes two sectors: a homogeneous goods sector that stabilizes aggregate costs, and a differentiated, polluting sector responsible for emissions. Firms in the differentiated industry sell their varieties into aggregate demand in both domestic and foreign markets. All equations are conditioned on the presence of the homogeneous sector.[9]

The model differs from the stylized version in Section 2 in two main ways, beyond explicitly specifying the functional forms. First, it explicitly incorporates the general equilibrium conditions faced by firms under the ETS policy. Specifically, firms are influenced on the output side by aggregate prices and foreign trade, which induces an indirect form of carbon leakage. On the input side, the carbon cap - which gets increasingly stringent over time - determines the marginal price of carbon through the budget clearing mechanism. For readability, aggregate variables are denoted by uppercase letters, while firm-level variables are represented in lowercase.

Second, the stylized model treats the clean technology adoption as a purely discrete decision. I generalize this by introducing multiple discrete adoption types, varying from minor to substantial (i.e., $\chi \in \{1.05, 1.10, 1.20\}$). For tractability, I maintain the discrete nature of the decision, enabling it to be estimated using a discrete choice model.

[9] In the EU, the polluting sector ($\approx$ ETS) contributes roughly 40% of total emissions but represents only 12% of value added and 6% of employment. This disparity supports the inclusion of a non-polluting sector.

4.1.1 Firms

Production Firms are heterogeneous in their productivity z and their level of clean technology a. Productivity consists of a fixed component z^F and a transitory component z^T, such that $z = z^F + z_t^T$, following Bartelsman et al. (2013) and Sterk et al. (2021), among others. The transitory component z_t^T is assumed to be an independently and identically distributed (i.i.d.) random variable: $z_t^T \sim \mathcal{N}(0, \sigma_z^2)$. Here, σ_z denotes the standard deviation of idiosyncratic productivity shocks. The transitory component is independently distributed across firms and time, and independent of other state variables. Since the profit maximization problem is static, I drop the time index to improve readability. [10]

The firms operate given a constant-returns-to-scale production function $q = zl$ with technology z and the only input labor l. As a byproduct, emissions are modeled as the production of an additional good proportional to output, $e = \frac{q}{a}$. As a consequence, one could rewrite emission as an input of production (Copeland and Taylor, 2004; Shapiro and Walker, 2018). The abatement factor stands orthogonal to the production factor. It could be restated as a Leontief production function, $q = \min[zl, ae]$ in the short term, where the degree of clean technology a acts as the share of e in the production function. By incorporating emissions in this manner, this approach diverges from the literature, which often assumes freely adjustable elasticities of substitution between emissions and other inputs such as Cobb-Douglas (Copeland and Taylor, 2004; Shapiro and Walker, 2018) or aims to estimate the elasticities in a static CES production function (Farrokhi and Lashkaripour, 2022). The dynamic nature of this model - implied by the sunk abatement decision - creates a long-term elasticity of substitution between the two factors of production - emissions and the input factor. One could interpret this as a putty-clay model where the vintage gets adjusted directly rather than the underlying capital stock (Atkeson and Kehoe, 1999).

A major benefit of considering it orthogonal to productivity lies in its interpretation where the taxation is a deviation from the BGP and the only way to return to the BGP at the firm level is for the abatement technology a to reach infinity, considering it a backstop technology (Golosov et al., 2014). This way of writing would also be in line with the literature on the misallocation of resources (Hsieh and Klenow, 2009; Restuccia and Rogerson, 2008) where the cost of emissions can be seen as the firm-specific taxation ($z\zeta/a$) relative to the optimal output of the firm.

Preferences and demand Demand in the home country for the domestically produced goods is a downward sloping function of the domestic price:

$$q_{hd} = p_{hd}^{-\eta} P_h^{\eta} Q_h, \tag{3}$$

[10] In the baseline model, I exclude an AR(1) process in z^T to avoid the computational burden of an additional state variable. However, in the appendix, I include an extension where the persistence parameter is estimated separately and rerun the model. The results remain largely unchanged.

where Q_h and P_h denote the composite aggregate demand for the home country and the associated price index respectively. η indicates the output elasticity. The demand for the exported good depends on the aggregate foreign price $q_{hx} = p_{hx}^{-\eta} P_f^{\eta} Q_f$, with, when countries are symmetric, $p_{hx} = \tau p_{hd}$, where $\tau > 1$ represents the iceberg cost.

Profit maximization Demand in the home and foreign country is separable. Given the demand for each variety in the home country, combined with the production function and externality, it gives the domestic d profit function for firms operating in the home country h:

$$\pi_{hd} = \max_{p_{hd}} p_{hd} q_{hd}(p_{hd}) - W l_{hd}(p_{hd}) - W \zeta e_{hd}(p_{hd}), \tag{4}$$

where W is the aggregate wage, normalized to 1. Given the demand for both input factors, the profit maximization w.r.t. price yields the following price setting:

$$p_{hd} = \frac{\eta}{\eta - 1} \frac{W c_h}{z}, \tag{5}$$

where the first term captures the constant markup and the second the marginal cost of the firm. The latter depends on the underlying productivity z, aggregate wage W and a cost function $c_h \equiv (1 + z\zeta/a)$. Quick comparative statics indicate that as $a \to \infty$ then $c \to 1$. Intuitively, the policy has no direct effect for extremely clean firms. One can interpret this limit-case as the implementation of a firm-specific backstop technology (Golosov et al., 2014) where its individual emissions are zero. Likewise, the degree of misallocation in a pure economic sense (Hsieh and Klenow, 2009; Restuccia and Rogerson, 2008) is non-existent when the price of carbon is zero ($\zeta \to 0$). Firms do not internalize this externality outside of the Pigouvian tax.

Profits generated from exports π_{hx} are derived in a similar manner as the price of the exported equals $p_{hx} = \tau p_{hd}$. I do not include a fixed cost of exporting because firms in the EU ETS are typically large and active in internationally traded sectors, making them likely to export.

Total profits are $\pi_h = \pi_{hd} + \pi_{hx} + \kappa \tilde{z}$. Here κ represents the lump-sum subsidies provided by the home government to domestic firms to mitigate carbon leakage. These subsidies are allocated based on productivity shares $\tilde{z}$, resembling the grandfathering approach used by the EU for distributing emission certificates to larger firms (Martin et al., 2014).[11]

Firms in the foreign country have the same profit function π_f, absent of the two policies ($\kappa, \zeta = 0$).

Technology adoption To better reflect realistic decision-making, I relax the discrete nature of the clean technology decision to one with various ranges of the size of the investment decision: $\chi = \{1.05, 1.10, 1.20\}$. One can interpret this differential as ranging from incremental to disruptive. Formally, the cost of the

[11] Note that in the baseline model, there is no exit decision. The share of certificates does not directly affect outcomes and primarily serves to satisfy the budget-clearing condition.

investment decision:

$$k(\chi) = k^F \mathbb{1}\{d_t \neq 0\} + k^C(\chi - 1)^2 \mathbb{1}\{d_t \neq 0\}, \tag{6}$$

where $k^F, k^C > 0$. The investment cost $k(\chi)$ is unscaled and it is composed of a fixed cost of investing k^F and a convex adjustment cost k^C expressed as a function of the percent investment rate χ (Cooper and Haltiwanger, 2006). Firms only pay the cost if they decide to adopt ($d_t \neq 0$). The first component governs the discrete decision of adopting a novel technology, while the latter restricts the degree of the step-size increase in the adoption decision.

4.1.2 Aggregate distribution and Market clearing

Number and distribution of firms The measure of firms active in the economy is normalized and can be stated as the sum over the firms in the two countries $N_t = N_{h,t} + N_{f,t} = 1$. I introduce Υ_t as a summary statistic for all aggregate components and reintroduce the time dimension t for clarification. I assume symmetry in the distribution at time t=0: $N_{h,0} = N_{f,0} = \int \mu_f(z_0, a_0)$. The distribution of firms in the home country exhibits the following motion:

$$\mu_h(z_{t+1}, a_{t+1}) = \int_{z*a} Pr(d_t|a_t, z_t, \Upsilon_t) F(a_{t+1}, z_{t+1}|z_t, a_t, d_t) d\mu_h(z_t, a_t), \tag{7}$$

where the transition operator $F(a_{t+1}, z_{t+1}|z_t, a_t, d_t)$ describes the transition law and encompasses both the endogenous effect of the discrete choice to d_t as well as the new draw of the idiosyncratic productivity component z^T. It is time-invariant and governs the stochastic evolution of the states given current choices. The probability function $Pr(d_t|a_t, z_t, \Upsilon_t)$ denotes the probability of each decision. Note that since the model is in a state of transition, the decision operator d_t depends on time t as does the distribution of firms.

The evolution of the distribution of firms in the foreign country only depends on the evolution of the productivity parameter:

$$\mu_f(z_{t+1}, a_{t+1}) = \int_{z*a} F(z_{t+1}|z_t) d\mu_f(z_t, a_t), \tag{8}$$

where the transition law only depends on the underlying productivity z_t.

Output market clearing The clearing conditions in this model are threefold: the budget constraint of the home government and the clearing of the output market. Since these conditions are considered within a single period, the time index t is omitted for clarity and readability. Aggregate demand allocated to the differentiated good in both countries is stable and normalized to $I_h = Q_{h,t} P_{h,t} = 1$ and $I_{f,t} = Q_{f,t} P_{f,t} = 1$. The home demand satisfies:

$$Q_{h,t} = \left[\int_{z*a} q_{hd,t}(z_t, a_t)^{\frac{\eta-1}{\eta}} d\mu_h(z_t, a_t) + \int_{z*a} q_{fx,t}(z_t)^{\frac{\eta-1}{\eta}} d\mu_f(z_t, a_t) \right]^{\frac{\eta}{\eta-1}}, \tag{9}$$

with domestically produced $q_{hd,t}(z_t, a_t)$ and imported goods $q_{fx,t}(z_t)$. Note that the latter does not depend on the level of clean technology, as the foreign price of carbon is zero.

The associated price index for the home country can be summarized as:

$$P_{h,t} = \left[\int_{z*a} p_{hd,t}(z_t, a_t)^{1-\eta} d\mu_h(z_t, a_t) + \int_{z*a} p_{fx,t}(z_t)^{1-\eta} d\mu_f(z_t, a_t) \right]^{\frac{1}{1-\eta}} \tag{10}$$

The aggregate demand and the price index for the foreign market are set equally.

Carbon market clearing The carbon policy is modeled after the EU ETS. One major implication is that the government sets the supply of emissions, $E^S_{h,t}$ with a cap, and leaves it up to market forces to determine the price of carbon. Firms are price-takers in the input market for carbon certificates. The optimal price of carbon at time t is determined by:

$$\zeta^*_t = \arg\min_{\zeta_t} \left| E^S_{h,t} - E^D_{h,t}(\zeta_t) \right|, \tag{11}$$

where the aggregation of emissions follows: $E^D_{h,t} = \int_{z*a} e(z_t, a_t) d\mu(z_t, a_t)$. The idiosyncratic demand for carbon at time t depends on the price of carbon as well as the level of the state variables (z_t, a_t).

Government budget clearing Finally, the budget clearing of the home government entails:

$$\int_{z*a} \zeta e_t(z_t, a_t) d\mu_h(z_t, a_t) = \int_{z*a} \tilde{z}_t \kappa_t d\mu_h(z_t, a_t), \tag{12}$$

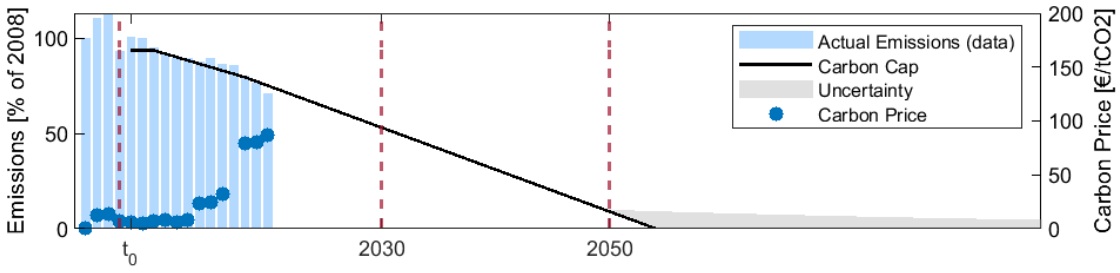
where the left-hand side denotes the government revenue through the taxation of the emission output. The right-hand side redistributes this revenue as a lump-sum subsidy, κ_t, to all domestic firms based on an allocation share determined by the ratio of the persistent productivity component of each firm ($\tilde{z} = (z^F)^\eta / \int_j (z^F_j)^\eta dj$). The government's budget constraint determines the aggregate value of κ_t. Since κ_t enters profits as a lump-sum transfer, it does not affect marginal outcome of the firms.

4.1.3 Green transition

The economy transitions from an initial balanced growth path ($t = 0$), where firms do not internalize emissions and therefore do not deviate from their current situation, to a second state at time T, characterized by zero emissions across all firm types due to the implementation of the carbon policy and universal adoption of a backstop technology. The initial phase corresponds with Phase I of the ETS, which was used as a trial period with monitored but untaxed emissions (Cantillon and Slechten, 2023). The second phase results from technological progress that makes clean technologies widely accessible. To encourage firms to adopt these technology, the government sets a cap on aggregate emission output.

Technological progress In the baseline version of the model, the price of technology diminishes at an exogenous rate. I assume that the cost of the clean technology adoption evolves at a constant rate γ. The

Figure 3: Evolution of the Carbon cap



Note: Evolution of the European Union's carbon emissions. The black line represents the EU-imposed carbon cap, decreasing linearly by 2.2% of 2005 emissions each year after 2021 (and by 1.76% annually from 2013 to 2020). The blue bars shows the actual emissions of the manufacturing firms in this sample, normalized to 100 in 2008 and the blue dots the corresponding price of carbon for the ETS. The grey shaded area after 2050 implies uncertainty surrounding the ETS-endgame (Pahle et al., 2025).

evolution of the cost parameter can be written as:

$$k^j_{t+1} = \gamma k^j_t \qquad s.t.\ \gamma < 1 \qquad for\ \ j = \{Q, F\} \tag{13}$$

The progress rate covers both the convex (k^Q_t) and the fixed cost (k^F_t) of adoption at the same rate. The only constraint in the estimation is the fact that the evolution of the technology parameter has to be lower than 1.[12]

Aggregate technological progress evolves according to $z_{t+1} = \omega z_t$ with $\omega \geq 1$.

Carbon market The model is estimated based on t *Cap-and-Trade* of the EU, the EU ETS. The emissions cap follows a predetermined reduction path designed to achieve the benchmark target by 2050. This cap is set ex-ante (during Phase I) for the entire time horizon $t = 1, 2, \ldots, T$, and is known to all participating firms. Figure 3 illustrates the evolution of the carbon cap established by the European Union, as detailed in Section 3.

As a simplified solution to address potential failures surrounding the ETS endgame, I introduce a slight deviation from the original emissions cap. Starting in 2050, domestic emissions in the EU are expected to exceed the cap potentially due to a combination of financial frictions, coordination failures as well as the usage of banked certificates from earlier periods (Pahle et al., 2025). From that point onward, I assume that emissions decline at a yearly rate of 1% until full decarbonization is reached. Ignoring this adjustment would result in a sharp increase in certificate prices.

[12] Aside from being true to general economic progress, it is a necessary requirement for the convergence towards a decarbonized economy in T years. In the baseline model, γ is independent of demand. In Section 5, I relax this assumption and allow it to evolve in a Romer-style framework.

4.2 Estimation strategy

The estimation procedure stacks the estimation of the two state variables, their initial distribution, idiosyncratic decisions, and the iceberg trade cost outside the model. The model is then used to estimate the cost of either decision as well as the technological progress parameter of the adoption decision. Finally, the discount factor, $\beta = 0.92$, the elasticity of substitution, $\eta = 6$, and the yearly rate of technological progress, $\omega = 1.01$ are set a priori.

To avoid confusion, in this empirical section, I reintroduce the firm-level index i which was omitted in the earlier sections for readability. Unless stated otherwise, all variables retain the same definitions as before.

4.2.1 Outside the model

I begin by deriving the two state variables that characterize the distribution of firms: productivity ($z_{i,t}$) and clean technology adoption ($a_{i,t}$). Firm-level productivity is first obtained and decomposed into its persistent and transitory components. Using this foundation, I characterize its joint non-parametric distribution with productivity and construct a firm-level measure of clean technology adoption. Finally, I derive the trade cost τ.

Firm Productivity To estimate firm productivity, I assume firms operate under a value-added production function that is Leontief in emissions $e_{i,t}$ and their level of clean technology: $q_{i,t} = \min\{z_{i,t} f(k_{i,t}, l_{i,t}), a_{i,t} e_{i,t}\}$. The productivity term $z_{i,t}$ is obtained using a factor-share decomposition following Asker et al. (2014) where output is divided by a markup of 1.2 given the demand elasticity $\eta = 6$. Full details of this procedure are provided in Appendix D. The structural model in the paper remains agnostic about the exact composition of inputs within $f(.)$. To maintain focus and tractability, I treat these as a single composite input represented as labor.[13] This specification parametrically identifies the distribution of the transitory component $z^T_{i,t}$; the estimation procedure does not require any assumptions on the parametric distribution of the fixed component z^F_i.

Clean Technology The level of clean technology can be backed out through the Leontief assumption on the production function as well: $q_{i,t} = a_{i,t} e_{i,t}$. Note that this assumption on the production function assumes no specific complementarities of the emissions with either input. In contrast, it only requires a convex combination of inputs that generates this result of $a_{i,t}$.

Joint Distribution The joint distribution of clean technology and the persistent component of productivity is identified non-parametrically using a gridded approach where I bin both values into equally spaced

[13] This paper does not aim to speak to allocative changes in the input factor market through the adoption channel. Papers that look at the interaction between climate policy, technology adoption and factor allocation are Finkelstein Shapiro and Metcalf (2023) for employment and Lanteri and Rampini (2025) and Capelle et al. for capital. I consider the level of abatement technology $a_{i,t}$ as a convex combination of the input mix where the composition of variables does not matter.

grids. More specifically, each bin accounts for a step-size increase of 5% in clean technology (see Appendix D). This yields 165 grids for a and 5 for z. Quick summary statistics show that, indeed, more productive firms have higher levels of clean technology independent of the sector they are in (Copeland et al., 2022; Shapiro and Walker, 2018), yet the relationship is not linear.

Adoption Decision The adoption decision itself is not directly observed in the data; instead, what is observed are changes in the clean technology variable $a_{i,t}$. Under certain assumptions, these observed shifts in $a_{i,t}$ can be interpreted as the structural adoption of new and improved technologies. Due to its irreversible nature, investment decisions tend to be large-scale and infrequent (Cooper and Haltiwanger, 2006; Caballero et al., 1997). I develop an approach, inspired by Bessen et al. (2019), to infer the timing and magnitude of technology adoption from observed changes in $a_{i,t}$. This method discretizes the often-used convex adjustment cost framework commonly used in models of firm dynamics (Asker et al., 2014; Clementi and Palazzo, 2016; Cooper and Haltiwanger, 2006; David and Venkateswaran, 2019, ...) and applies it to the clean technology margin. Formally:

$$d_t(\chi) = \mathbb{1}\underbrace{\{\chi a_{i,t+1} \geq a_{i,t}\}}_{\text{Transitory state}} \underbrace{\left\{\chi \frac{1}{B}\sum_{b=1}^{B=2} a_{i,t+b} \geq \frac{1}{W+1}\sum_{w=0}^{W=2} a_{i,t-w}\right\}}_{\text{structural state}}, \tag{14}$$

where $d_t(\chi)$ indicates the decision to invest for various cutoff factors $\chi \in \{1.05, 1.10, 1.20\}$. The indicator is made up out of 2 states: The *transitory state* governs the lumpiness of the increase in clean technology. Adoption decisions do not occur often. The second *structural state* governs the long-term absorption of clean technology in a new, structurally higher environment. Following the literature (Cooper and Haltiwanger, 2006), I consider the decision to invest to occur one year prior to the actual outcome change.

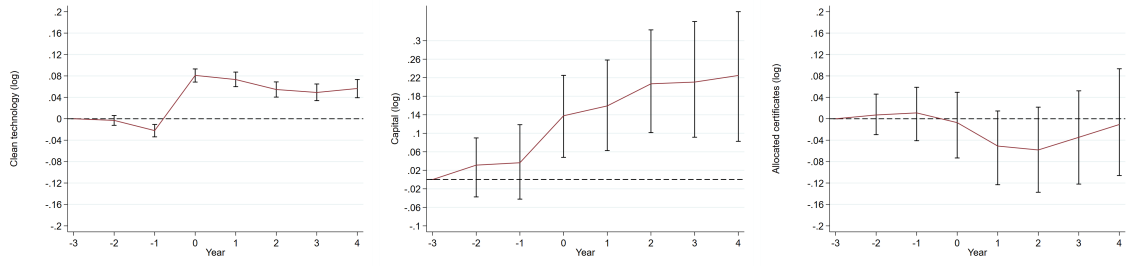
To confirm that this genuinely reflects investment, I carry out several robustness analyses. First, I apply a falsification exercise where I test the co-movement of capital and labor in an event-study design (Sun and Abraham, 2021), using the difference-in-differences notation following Callaway et al. (2021). Emission reductions occur at staggered times, with the treatment defined at $\tau = 0$ and decision at $\tau = -1$. Firms that reduce emissions once ($D_i = 1$ if $\sum_t d_{i,t} = 1$) are compared to those that never do ($D_i = 0$ if $\sum_t d_{i,t} = 0$), excluding repeating investors.[14] The specification is:

$$ln(x_{i,t}) = \sum_{\tau=-3}^{T=4} \alpha_\tau G_\tau D_i + \iota_{t,c} + \gamma_{t,s} + \delta_i + \varepsilon_{i,t}, \tag{15}$$

where $x_{i,t}$ includes material cost, capital, and emissions; G_τ indicates years relative to the treatment; ι_t, δ_i are time and firm fixed effects. Errors are clustered at the firm level. Following Young (2019), the control group is a random sample of firms that never adopt, which are given placebo treatments. This process is bootstrapped 1000 times. Results in Figure 4 show the median over the horizon with 5 - 95% confidence

[14] To include repeating investors, a first-difference regression analysis is conducted in Appendix D. The results stay consistent.

Figure 4: Co-movement variables during 'Technology adoption' with $\chi = 1.10$



(a) Clean technology (b) Tangible fixed capital (c) Allocated certificates

Note: Co-movement of variables with the increase in abatement technology $a_{i,t}$. The adoption decision $d_{i,0} = 1$ occurs in the year zero. The timing is normalized to the year -3 as this is the baseline for which the technology adoption decision moments start to be binding. Confidence bands are set at the 5 to 95% interval.

bands. The event window extends one year beyond the reduction to test for mean reversion. The estimated growth in capital is statistically significant and different from zero, indicating that an increase in technology adoption of about 10% is accompanied by a proportional rise in capital.[15] The effect on the allocation of emission certificates remains statistically indistinguishable from zero, suggesting that the observed effects are not driven by the closure or dismantling of emission-intensive output.

Second, to verify that the firms investing in cleaner technologies are also those experiencing greater emission reductions, as discussed in Section 3, I regress an indicator for whether a firm invests at least once during a given phase of the ETS on its ex-ante characteristics.

Third, to assess whether the ETS provides an independent incentive for such investment rather than the observed behavior being driven solely by other factors such as oil prices or shifts in aggregate demand, I regress the investment indicator on the interaction between the carbon price and the firm's level of clean technology, thereby identifying a firm-level responsiveness to carbon prices for less abatement intensive firms. Detailed regression specifications and supplementary figures are provided in Appendix D. Overall, the results indicate that investment decisions are concentrated among a common subset of firms and are more frequent when carbon prices are high. [16]

This outcome-based approach offers distinct advantages relative to using input-based approaches: it avoids the need to classify specific firm actions as abatement and sidesteps the difficulty of estimating the output additionalities of such efforts. What can be considered a 'clean' investment for some firms can as well be considered a polluting one for firms operating in different locations/industries or with different

[15] This pattern is indicative of the non-monotonic relationship between capital and emissions. While the literature typically finds capital to be highly correlated with emissions, increases in capital can also be associated with emission reductions.

[16] All estimates for Phase III are statistically significant. The results for Phase II are more suggestive, with several coefficients not attaining conventional levels of significance. This weaker statistical precision likely reflects (i) the distortionary effects of the financial and subsequent sovereign debt crises in the European Union during this period, (ii) unobserved heterogeneity across firms or sectors, and (iii) the presence of 'low-hanging fruit' dynamics that the model cannot fully capture.

production functions/technologies. Moreover, this approach aligns with the approach taken by the EU in designing the ETS, where the aim was to give firms maximal autonomy in choosing how to achieve reductions in emission intensity. As a final remark, this methodology has a one-to-one mapping to the quality-ladder framework in the structural model.

Foreign trade During the first phase of the EU ETS (2005 - 2007), the policy had minimal distortional impact on the economy (Cantillon and Slechten, 2023; Colmer et al., 2020), effectively rendering the domestic economy symmetric to the foreign country at this time. This renders the estimation of the trade cost to be modeled using a standard trade framework (Arkolakis et al., 2012). The cross-border trade cost is quantified using a Head-Ries index (Head and Ries, 2001; Antras and Chor, 2018).

To obtain τ, I utilize EUROSTAT data on the ratio of EU industrial exports to the rest of the world relative to total sales for the years 2005 - 2007. This approach assumes that the trade behavior of ETS-covered firms is similar to the overall aggregate statistics and that the trade would have remained relatively constant absent of the carbon policy. This assumption is supported by the significant size and influence of these large firms within total exports.

4.2.2 Estimation of the model

I initiate the estimation strategy by revising the discrete choice function of Equation 1 in Section 2 as a function of the idiosyncratic agents. The payoff of agents in period t depends on their observed state variables a_t and z_t, the exogenous aggregate economic conditions Υ_t, and an agent-specific variable $\varepsilon(d_t)$, not observed by the econometrician. The general equilibrium version of the ex ante value function from Equation 1 can be written as $V_t(a_t, z_t|\Upsilon_t)$, where Υ_t is a vector that encapsulates the exogenous government policies and the endogenously determined general equilibrium prices.[17] This value function can then be expressed as an average over all possible realizations of the unobserved shock ε_t.

$$V_t(a_t, z_t|\Upsilon_t) \equiv \int \tilde{V}_t(a_t, z_t, \varepsilon_t|\Upsilon_t) g(\varepsilon_t) d\varepsilon_t, \tag{16}$$

where ε_t is assumed to be additively separable and distributed according to an extreme value distribution $g(\varepsilon_t)$.

The solution method is based on backward induction. The horizon of the dynamic model is infinite and non-stationary. Since the parameter of clean technological progress $\gamma < 1$ ensures that the model transitions towards a balanced growth environment where none of the firms finds it optimal to deviate, as all firms eventually reach a level of $a \to \infty$. In practice, I follow Acemoglu et al. (2016) and consider that full decarbonization is reached in at least 1000 years from now. The optimal decision at time $T \leq 1000$ for all agents

[17] As Υ_t is partially unobserved, it has to be solved through the model. It is a function of the full distribution of firms.

will be not to deviate from the current situation:

$$\arg\max_{d_T}[v_T(z_T, a_T, d_T|\Upsilon_T) + \varepsilon_T] = 0 \qquad \forall z_T, a_T \tag{17}$$

From T, backward induction will be applied to solve the model. The ex ante value at time $T-1$ is a function of all potential probabilities and the error term ε_T:

$$V_{T-1}(z_{T-1}, a_{T-1}|\Upsilon_{T-1}) = \int \sum_{d_{T-1}} [v_T(z_{T-1}, a_{T-1}, d_{T-1}|\Upsilon_{T-1}) + \varepsilon_{T-1}]g(\varepsilon_{T-1})d\varepsilon_{T-1} \tag{18}$$

I estimate the model over the full state space in order to account for the aggregate equilibrium effects that influence the adoption decisions. The probability of observing an arbitrary choice d_t, conditional on the value function v_t, is given by:

$$Pr_t(d_t|a_t, z_t, \Upsilon_t) = \frac{exp(v_t(a_t, z_t, d_t|\Upsilon_t))}{\sum_{d_{jt}} exp(v_t(a_t, z_t, d_{jt}|\Upsilon_t))}, \tag{19}$$

where j governs the discrete choices - do nothing or invest with λ percent. The probability of all options j yields $\sum_{d_{j,t}} p_t(d_{j,t}|a_t, z_t, \Upsilon_t) = 1$ and $Pr_t(d_{j,t}|a_t, z_t, \Upsilon_t) \geq 0 \; \forall \; j$. Forming these probabilities for each individual i at each time period t yields the components necessary to set up a likelihood function to obtain the parameter vector $\hat{\theta} = \{k_0^F, k_0^Q, \gamma\}$ which is a function of the three unknown variables governing the adoption decision.

While backwards recursion is done until time T is reached, the data is only available until the year 2018 - stated as 11 years from the steady state in phase I when phase II starts (2008, where $t = 0$). With N_h firms, the estimated parameters are obtained by:

$$L(\theta) = \sum_{t=1}^{10} \sum_{i=1}^{N_h} ln[Pr_t(d_{i,t}|a_{i,t}, z_{i,t}, \Upsilon_t, \theta)] + ln[F(a_{i,t+1}, z_{i,t+1}|a_{i,t}, z_{i,t}, \Upsilon_t)], \tag{20}$$

where $F(a_{i,t+1}, z_{i,t+1}|a_{i,t}, z_{i,t}, \Upsilon_t)$ is the transition operator associated with the probabilities of the discrete choices. Note that this operator is the same as in the aggregate distribution evolution as described in Equation 7. From here, one can obtain $\hat{\theta}$ that maximizes the likelihood function.

5. Results

This section presents the results of the model estimation for the EU carbon policy. It first establishes the baseline scenario with key parameter estimates. Then I delve into how the underlying mechanism shapes distributional implications of the carbon policy - this also acts as a validation exercise. Next, the analysis decomposes the trajectories of aggregate emissions and domestic sales, tracing their evolution to two main drivers: the direct impact of carbon policy on marginal costs and idiosyncratic decisions. Third, I examine the role of different policy instruments such as a CBAM and adoption subsidies. Finally, it explores

sensitivity checks and applies an extension to clean technological change.

5.1 Baseline results

The parameter estimates are reported in Table 2. The table includes both point estimates and their associated standard errors of the costs associated with the adoption decision. The results reveal a notable disparity between the two types of costs: the estimated evolution parameter γ indicates that the sunk cost of adoption declines at an average rate of 1.445% per year (=1-0.9855), reflecting a gradual progress of clean technology.

Table 2: Parameter estimates

Variable:		Value:	s.e.:
Fixed component Abatement adoption cost	(k_0^F)	0.6469	(0.0547)
Convex component Abatement adoption cost	(k_0^Q)	2.412	(0.908)
Evolution abatement cost	(γ)	0.9855	(0.0033)

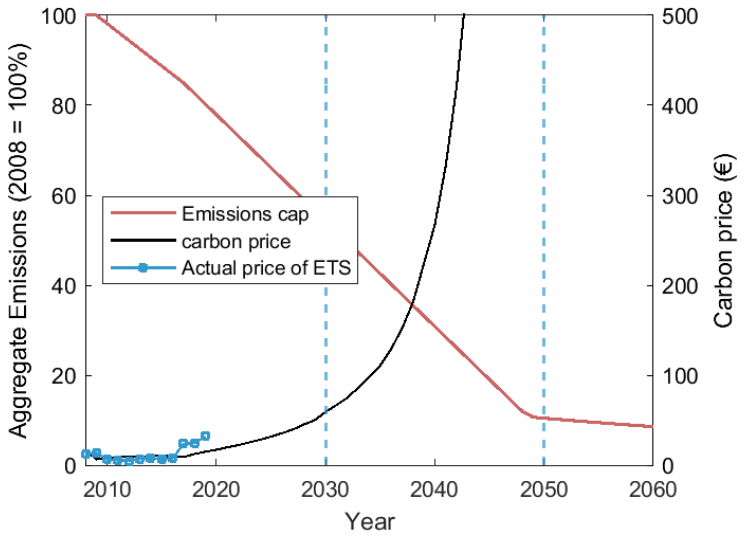
Note: The table presents the estimated results for the parameters in the log likelihood function: The fixed and convex component of the adoption cost and their technological progress. Standard errors are documented in parenthesis in the second column.

As an initial result and a validation step, Figure 5 presents the price of carbon ζ_t for which the emission market clears (Equation 11). The difference between the implied price and the actual price is largely the same even when abstracting from aggregate uncertainty and permit savings, or excluding firms outside the manufacturing sector. Similar to observed spot carbon prices, the market-clearing price starts extremely low. The model predicts that the price increases at an exponential rate until 2050, and then gradually declines toward zero as all firms transition to their backstop technology. As time progressing the price has less of an impact on aggregate output through the decoupling of emissions and output.

5.1.1 Firm Inequality

The model provides outcomes that highlight the role of heterogeneity in reducing aggregate emissions. As in the data, large polluting firms are more likely to adopt novel technologies and thus drive the aggregate result. Figure 6 shows the cumulative annual probability of adopting new clean technologies for the 3 different values of $\chi = \{1.05, 1.10, 1.20\}$ along the productivity and the clean technology dimension for firms operating in the home country for the years 2013 to 2018 - mimicking phase III of the policy. In these 5 years, approximately 15% of all firms have adopted at least once an improved technology of 5%. The figure demonstrates that the probability of engaging in either decision is increasing in productivity and decreasing in clean technology. The rate is substantially larger for the larger steps on the quality ladder. Additionally,

Figure 5: Carbon price and market clearing



Note: Evolution of domestic emissions (red) and its corresponding price that ensures market clearing (Equation 11) in black. In blue the actual spot price of the EU ETS is depicted for the years 2008 to 2020 (i.e. before the 21 policy change (Pahle et al., 2025)) the years 2030 and 2050 show the *fit-for-55* and *net zero* objectives by the EU.

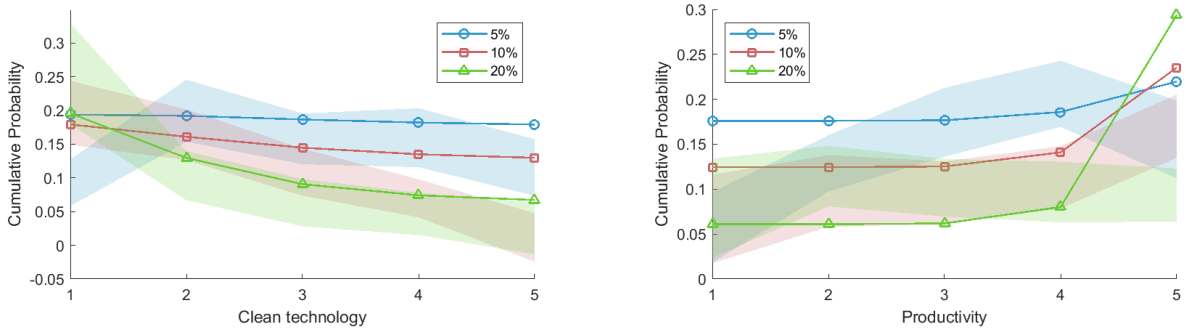
with the exception of the overestimation of the firms in the most productive quintile that adopt at a sharp rate of 20%, these results show that the model provides a good fit for the actual data.

I then turn towards the outcome variable. More specifically, I decompose emissions along these margins to see how these adoptions drive firm-level outcomes. Individual emissions can be decomposed into two distinct channels. The first is the direct channel, whereby carbon taxation induces a reduction in output through a rise in the marginal cost. The second is the indirect channel, corresponding to abatement adoption, through which firms adjust their emissions by adopting novel technologies. The evolution of emissions is a function of the current emissions minus the counterfactual emission output in the business-as-usual scenario: $\Delta e(a_t, z_{t+1}) = e_{t+1}(a_{t+1}, z_{t+1}) - e^*_{t+1}(a_t, z_{t+1})$. Decomposing individual growth of emissions conditional on keeping aggregate prices constant (given the continuous mass of firms assumption) leads to the two specific components:

$$\Delta e(a_t, z_{t+1}) = \underbrace{[e_{t+1}(a_t, z_{t+1}) - e^*_{t+1}(a_t, z_{t+1})]}_{\text{direct margin}} + \underbrace{[e_{t+1}(a_{t+1}, z_{t+1}) - e_{t+1}(a_t, z_{t+1})]}_{\text{Adoption}}, \tag{21}$$

where a_{t+1} denotes the potential future level of abatement technology. The first term captures the effect of the taxation relative to the business-as-usual scenario. This highlights the implications of the carbon policy absent any dynamic decisions. The evolution of the total price index is captured in the direct effect. The second term reflects the impact of adopting novel technology compared to not adopting. If the firm

Figure 6: Idiosyncratic decisions by dimension (2013 - 2018)



(a) Clean technology margin (a_t)

(b) Productivity margin (z_t)

Note: Cumulative probability of adoption between 2013 and 2018 along 2 dimensions of heterogeneity: the clean technology margin (left panel) and the productivity margin (right panel). They are considered as a Bernoulli trail: $1-\prod_{t=2013}^{2018}(1-d_t(\chi))$. The margins are equally spaced bins based on the ex-ante levels of clean technology and productivity (i.e. at the year 2013). The more right on the x-axis the higher average clean technology / productivity of the quintile. the colors show the adoption probabilities for the 3 different values of $\chi = \{1.05, 1.10, 1.20\}$. The shaded area represents the 95% confidence intervals in the data for their respective decisions given the equivalent ex-ante value of the margin, conditional on the average level of the other margin. The results for phase II of the ETS can be found in Appendix E; they align as well with the data and are in relatively same magnitude.

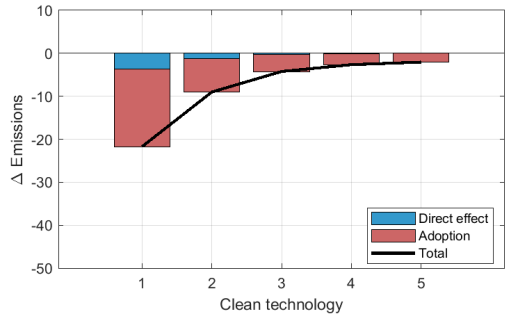
chooses not to adopt, the second term cancels out, leaving only the direct effect. While on average, the direct margin is negative, if the level of clean technology is sufficiently high, the direct effect might yield positive results. This highlights the reallocation effect through the general equilibrium prices, where greener firms find it easier to produce when prices are high in a setting where quantities and emissions are tightly linked. The implications of the adoption channel are always negative.[18]

Output can be decomposed similarly with $q_h = q_{hd} + q_{hx}$. In contrast to emissions, the adoption channel mitigates the negative effect a carbon policy has on output. When firms have reached higher abatement levels, the implications of a policy are less severe on their output.

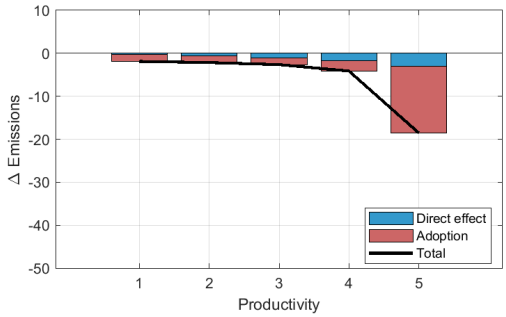
The results for both decompositions for the period of 2013 to 2018 are presented in Figure 7. Overall, the results show that polluting firms adopt more and therefore reduce emissions to a greater extent. On average, firms reduce their output due to the policy. Note that there is some reallocation to the clean firms as they can produce more output following the implications of the direct margin (Figure 7a and Figure 7c). However, the majority of the reductions are carried by the adoption channel. The adoption channel has a mitigating effect on output losses (Figure 7c). For productivity, the relationship is reversed:

[18] When adoption is zero, and given the assumption of a CES structure for aggregate demand, this implies that, conditional on their current level of clean technology, the effects of the carbon policy are identical across firms. However, more polluting firms adopt new technology at higher rates, resulting in larger emissions reductions for these firms. Variations in the direct effect across different margins in the figures can be attributed to distributional patterns, with less productive firms generally being more polluting.Although the results indicate that size-based inequality increases, this does not imply welfare losses due to the CES demand structure, which ensures efficient market outcomes.

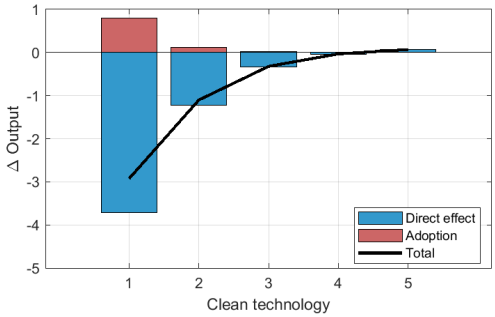
Figure 7: Idiosyncratic outcome variables by dimension (2013-2018)



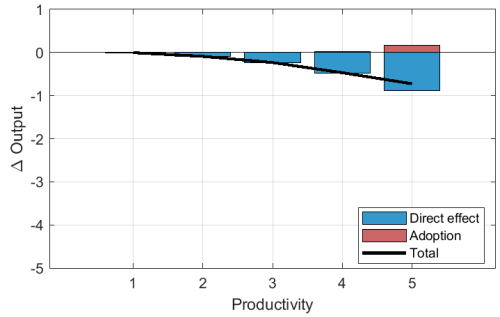
(a) Clean technology margin



(b) Productivity margin



(c) Clean technology margin



(d) Productivity margin

Note: Changes in emission output Δe and produced output Δqh (where $q_h = q_{hd} + q_{hx}$ for each firm type z and a) between 2013 and 2018 along 2 dimensions of heterogeneity: the clean technology margin (left panels) and the productivity margin (right panels). The changes Δx (for $x \in \{e, q\}$) are decomposed following Equation 21. The deviations are written in percentages from their situation in 2013. The blue component in the bar shows the implication of the direct effect of the taxation. In red, the effect of the adoption is given. The black line shows the total reduction. The relationship between the results and the data can be found in Appendix E.

more productive firms both adopt more and achieve larger emissions reductions.

5.1.2 Aggregate outcomes

The policy ensures that domestic emissions converge toward zero (as indicated earlier in Figure 3). The central question is how much of the observed evolution in aggregate demand can be attributed to the idiosyncratic abatement decisions. To answer this, I decompose the dynamics into two distinct components. The first is the direct margin, which isolates the effect of carbon policies ζ_t on marginal costs while holding the firm distribution fixed (i.e., abstracting from idiosyncratic decisions). This channel also captures the impact of the policy on carbon leakage through the general equilibrium effects. The second margin is the indirect effect, which arises from changes in the firm distribution induced by the adoption of new

technologies.[19] Formally:

$$\begin{aligned}\Delta Q_{h,t} &= Q_{h,t} - Q^*_{h,t} \\ \Delta Q_{h,t} &= \underbrace{[Q_{h,t}|_{d_t=0} - Q^*_{h,t}]}_{\text{direct margin}} + \underbrace{[Q_{h,t} - Q_{h,t}|_{d_t=0}]}_{\text{Adoption}},\end{aligned} \tag{22}$$

here, $Q^*_{h,t}$ denotes the counterfactual domestic demand under business-as-usual (BAU), i.e. in the absence of carbon policies. As $t \to \infty$, the adoption channel can mitigate the extent to which carbon policy increases the marginal costs and affects domestic demand. Due to the Leontief production function assumption in emissions, differences in aggregate outcomes can be interpreted as a direct measure of misallocation in the economy; affecting the economy negatively by default. The level of misallocation is a function of the carbon price ζ_t and a weighted average abatement level $\bar{a}_t$. As firms adopt novel technologies and $\bar{a}_t$ increases over time, the distortive impact of misallocation diminishes. In the long run, misallocation disappears and the economy converges to the BAU scenario, where the policy is absent and output grows at a constant rate. The key question is how much of this decline can be attributed to the direct channel versus the adoption channel at each point in time.

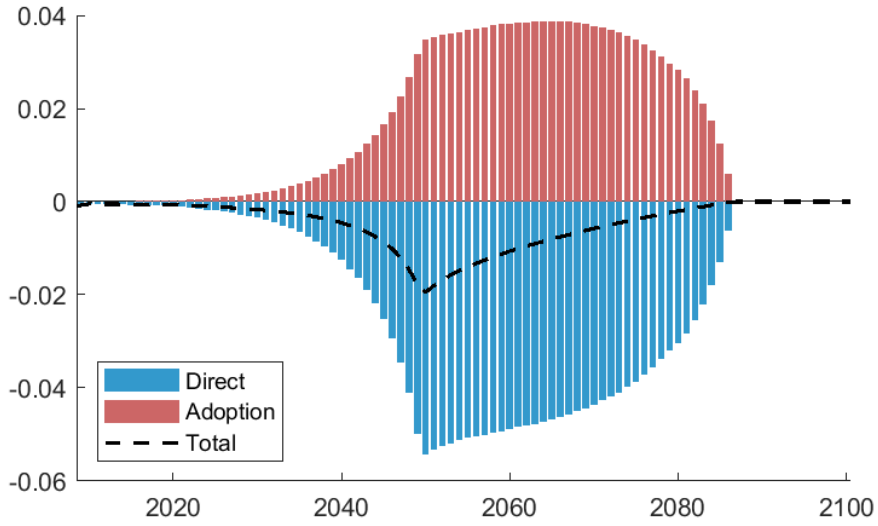
Figure 8a documents the evolution of aggregate domestic sales relative to the BAU scenario. The dashed line represents the total cost of the carbon policy on domestic demand. The cost of the policy is almost negligible in the first few years after the implementation, with only $\sim 0.4\%$ in 2025. The cost peaks at around $\sim 1.9\%$ in 2050. Appendix E provides additional results on the revenue decomposition as well as further decomposed first-order results: the cost are significantly higher for domestic sales with $\sim 1.0\%$ in 2025 and $\sim 5.0\%$ in 2050, as they are substituted away for non-exposed goods produced by the foreign market (i.e. carbon leakage). Further decompositions up to a first-order approximation demonstrate that the foreign market is significantly worsening the effectiveness of carbon policy on output losses.

After decomposing the total effect into either channel, one observes that the adoption channel has a substantial mitigating effect. Given the level of clean technology in 2008 (holding the distribution fixed), the current projected carbon price would result in a 6% reduction of domestic demand.[20] The ratio of the contribution of either channel is presented in Figure 8b. Where the contribution of the adoption channel increases incrementally to the year 2050, and then due to the loosening of the cap experiences a slight decline before increasing steadily to 100%. At 2025, approximately 40% of the reduction in demand is mitigated by adoption, while at 2050, this reaches almost 65%.

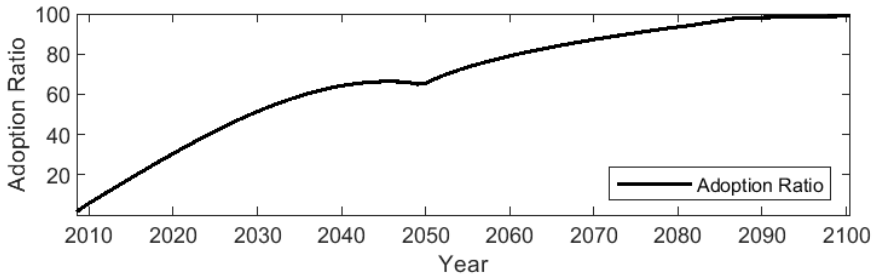
[19] Both effects can be further decomposed up to a first-order approximation. The decomposition of the direct effect would capture the static indirect implications of aggregate price movements - either foreign or domestic. Secondly, the decomposition of the indirect effect enables the visualization of the contribution of each technology adoption steps, parameterized by χ - i.e. are the results driven by the many incremental steps or by the few sharp adoptions. A full decomposition can be found in the subsection C.4 in the appendix.

[20] This is absent general equilibrium effects. Since output and emissions are tightly linked (via the Leontief assumption), eliminating domestic emissions would imply a proportional contraction in domestic production. Consequently, achieving net zero emissions would coincide with a net zero reduction in domestic output.

Figure 8: Decomposition aggregate demand (to 2100)



(a) Demand decomposition



(b) Ratio of adoption relative to direct component

Note: Figure 8a presents the decomposition of the evolution of aggregate domestic demand relative to the business-as-usual scenario. The dashed line shows the total effect of the carbon policy. It can be decomposed in two components. The direct effect of the carbon policy (blue) is calculated under a fixed distribution of firms and the adoption effect (red), which captures changes due to shifts in the firm distribution. Figure 8b provides the ratio of the adoption component over the direct component.

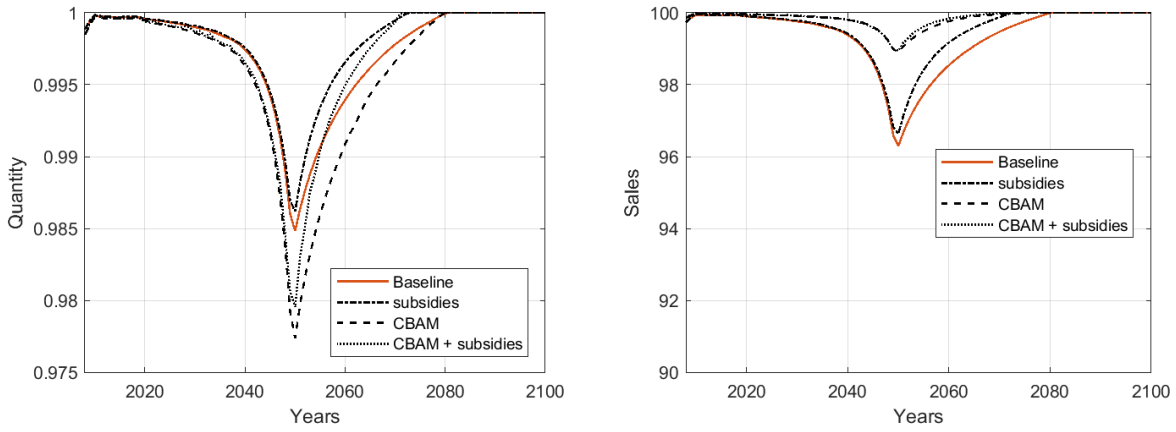
5.2 Policy instruments

This section begins with a counterfactual analysis motivated by the unilateral design of the EU ETS, which places domestic firms at a disadvantage relative to foreign producers not facing comparable carbon constraints. I connect to a large body in the trade literature that examines the design of a carbon border adjustment mechanism (Kortum and Weisbach, 2021; Farrokhi and Lashkaripour, 2022; Fowlie and Reguant, 2022; Fontagne and Schubert, 2023; Shapiro, 2021; Bohringer et al., 2017, ...), and extend it to a dynamic

framework along the transition path. [21] A second strand of discussion concerns the effective use of revenues generated by carbon policy - the recycling of the funds (Timilsina, 2022). Here, I examine their allocation in the form of subsidies aimed at facilitating abatement adoption.

Using the estimated parameters, I analyze a counterfactual scenario in which the domestic government had put a price on imported carbon emissions and another scenario in which the government used the income from the policy to subsidize abatement investments. I then explore the implications of either policy and their interaction and how they influence firm-level decisions and subsequently aggregate economic and environmental outcomes.

Figure 9: Outcome for two policy instruments



(a) Domestic demand

(b) Domestic sales

Note: Evolution of the outcome variables - aggregate domestic output (left panel) and domestic sales (right panel) for the two policy instruments: CBAM (dashed), adoption subsidies (dot dash) and their interaction (dotted) relative to the baseline model in red.

To compute the average rate of subsidies that the firm obtains, I start by calculating the total cost for either component. The ratio of income from the taxation relative to the total cost of adoption can be stated as:

$$s_t = \frac{\int_{z*a} \zeta_t e_t(z,a) d\mu_h(z,a)}{\int_{z*a} \sum_{\chi} k_t(\chi) pr(d_t(\chi) = 1 | z_t, a_t) d\mu_h(z,a)} \tag{23}$$

such that $k_t(\chi)' = \max\{(1-s_t), 0\} k_t(\chi)$ indicates the implied cost of adoption after discounting for the subsidy, where $'$ indexes the counterfactual nature of the variable. Note that this simplified subsidy schema is assumed to be static such that it has zero dynamic implications (i.e. the decreasing process of the cost $k^*_{1,t} = \gamma k^*_{1,t-1}$ does not depend on the subsidy scheme) and that the clean technology providers behave in perfect competition and do not reap any of the benefits of the subsidy. The subsidies are rebated in proportion

[21] A border tax adjustment is politically contentious (Timilsina, 2022). This framework does not allow for reactions from the foreign government to the implementation for such policy and thus considers a best-case scenario.

to all firms. In both scenarios, other variables remain equal to the baseline model.

The CBAM is modeled as a price of carbon on the imported goods to the domestic country: $p_{fx}(z_t, a_t)' = \tau\eta/(\eta-1)W/z_t(1+\zeta/a_t)$. Finally, I consider the scenario where both policy instruments are used and interact. Only domestic firms can make use of the subsidy scheme.

The results in Figure 9 illustrate the effects of both policy scenarios. Implementing a CBAM eliminates the comparative advantage of foreign competitors in the domestic market, but it also reduces overall domestic demand. While domestic firms gain a larger relative output share in the home market, their total sales remain below the business-as-usual level due to losses in export markets. This shift increases the implied demand for carbon, leading to a slightly higher carbon price and, in turn, a really modest increase in the adoption rate.

A rebate of subsidies increases the adoption rate and, as time continues, reduces the impact of the carbon policy on the aggregate outcome variable; both in the baseline as in the CBAM scenario. Additional results can be found in Appendix F.

5.3 Sensitivity analysis & Extensions

This section aims to test the sensitivity of the outcomes generated by the model to assumptions about key parameters. In particular, I examine the role of the exogenously set parameters β and η as well as the 2 parameters governing the trade flows, τ and foreign wage W_f. Finally, I consider the cutoff value $E_{h,T}$. As an extension, I treat the evolution parameter γ_t as endogenous and demand-dependent. Additional details on the construction, estimation, and calibration of these parameters can be found in Appendix F.

5.3.1 Sensitivity analysis

I begin by analyzing how the results respond to variations in the exogenously specified parameters: the discount factor β and the elasticity of substitution between goods η. A higher discount factor increases the adoption rate and reduces output losses. Because agents place more value on future gains, they are more willing to invest in adoption early. Similarly, a higher elasticity of substitution η slightly raises the adoption rate but also amplifies output losses. Greater substitutability intensifies competition across goods, which slightly boosts adoption incentives but increases output losses due to stronger price effects.

Next, I compare the baseline outcomes to scenarios with changes in the symmetric trade cost τ and the foreign wage W_f. An increase in τ has only a minor effect on demand and, consistent with the result that output is not primarily driven by reallocation, it leaves overall output largely unchanged. The same conclusion holds for variations W_f.

Finally, I examine the endgame - when emissions fall below 10% of their initial level - and test cutoffs of 0%, 5%, and 15%. The transition paths are similar, but the cutoff matters at the peak: driving emissions to zero nearly quintuples demand losses compared to stopping at 10%.

5.3.2 Endogenous technological progress

In the baseline version of the model, technological progress is modeled as an exogenous linear process in which the cost of adoption $k(\chi)$ decreases over time at a constant rate $\gamma = 0.9855$. I extend this baseline framework by constructing a Romer-type mechanism of technological progress. I consider a simple scenario of where the elasticity is equal to 1 (as in Acemoglu et al., 2012); thus it neither displays the ideas are harder to find pattern (Bloom et al., 2020) nor the building-upon giants effect (Jones, 2022). As before, all technological progress is reflected in the acquisition price of the new clean technology. The evolution of the cost of clean technological progress is given by: $\gamma_t = \left[1 + \psi \int_{z*a} \sum_{\chi} k_t(\chi) pr(d_t(\chi) = 1 | z_t, a_t) d\mu_h(z,a)\right]^{-1} \gamma_{t-1}$ where the parameter $\psi = 0.81$ denotes an composite form of research productivity. The model is calibrated so that the average endogenous rate of technological progress for the first 10 years replicates the exogenous baseline value of $1/10 \sum_{t=2008}^{2018} \gamma_t \approx 0.9855$. All other parameters remain fixed.

This setup yields a key contrast with models that view aggregate technological change as being primarily driven by innovation rather than adoption (Acemoglu et al., 2012; Hemous, 2016). Whereas my framework predicts that policy should optimally begin more gradually in order to smooth the welfare implications, the innovation-driven approach prescribes the opposite: ramping up policy aggressively to accelerate technological progress and expedite the transition to a green economy. This comparison here highlights the first-order magnitude of either.

The results in Appendix F indicate that allowing for an endogenous γ raises the aggregate adoption rate, which in turn lowers the carbon price and reduces the impact on aggregate output. The effects are moderate - not negligible, but not large either - with peak sales declining slightly from 4.5% to 4%.

6. Achieving net zero

This section examines the optimal domestic policy path that minimizes the cumulative loss in aggregate demand over the transition period. While the baseline analysis takes the evolution of the carbon cap as given (as presented in Figure 3), the model framework is well-equipped to evaluate alternative policy trajectories toward a low-carbon economy. Policymakers face a fundamental trade-off: faster decarbonization can come at the expense of economic output.

The optimal policy is benchmarked against a first-best scenario implemented by a global social planner, where both domestic and foreign regions coordinate efforts to reduce emissions. In this coordinated framework, the distortional effects of foreign competition are abstracted from. In the model, the domestic government pursues the EU's target of reducing emissions to 90% below 2005 levels by 2050, followed by a gradual convergence toward a backstop technology after time T. I then derive the optimal emissions path under both multilateral and unilateral policy frameworks and compare these with the current trajectory.

6.1 The government objective function

The home government seeks to maximize aggregate domestic demand (Farrokhi and Lashkaripour, 2022; Shapiro, 2021; Kortum and Weisbach, 2021), conditional on global emissions, by minimizing the sequence of the supply of aggregate domestic emissions $\{E^S_{h,t}\}^T_{t=t_0}$ until reaching net zero in 2050 (which implies $E^S_{h,T} = 0.10E_{h,0}$). Utility in the home country is defined as $U_{h,t} = Q_{h,t} - \lambda E_t$. Here the value of λ captures the negative aspect associated with total global emissions, $E_t = E_{h,t} + E_{f,t}$. The expression is intentionally specified in a simple linear form to facilitate interpretation of λ. Formally, the optimization problem for the home government can be expressed as:

$$\max_{\{E^S_{h,k}\}^T_{k=t_1}} \sum_{t=t_0}^{\infty} \beta_G^{t-t_0} \left\{ \left[Q_{h,t}(\{E^S_{h,k}\}^T_{k=t_1}) - \lambda E_t(E^S_{h,k}\}^T_{k=t_1}) \right] \right\} \tag{24}$$

where the sequence $\{E^S_{h,k}\}^T_{k=t_1}$ represents the emissions cap set by the domestic government for all periods k from t_1 to time T and publicly known to all agents. Although the government cannot take any further action after time T, it continues to discount future utility, which itself still depends on the previously established cap path and - I assume - an exogenous evolution of the path at an annual rate of 2% per year.

To make this problem numerically solvable, I impose additional structure on the path of the emissions cap: it depends solely on the initial stock of emissions E^S_{h,t_0} and a discount function $\theta(t)$:

$$E^S_{h,t} = \theta(t) E^S_{h,t_0} \tag{25}$$

where the discount function is specified as a second-order polynomial, $\theta(t) = b_1 t^2 + b_2 t + b_3$. As the discount function exhibits 2 degrees of freedom (i.e. at $t = 0$ and T the values are respectively 1 and 0.1), the system can be expressed in terms of a single unknown parameter, b_1. This formulation enables the function to express both concave ($b_1 > 0$) and convex ($b_1 < 0$) trajectories for the carbon path. To determine an optimal carbon policy, the parameters b_1 and the derived parameter $b_2(b_1)$ are solved simultaneously. Finally, I assume that the cap can never exceed a higher bound of ($E_{h,t} > E_{h,0}$) nor lower bound (i.e. $E_{h,t} < E_{h,T}$) for all $t \leq T$.

Rephrasing the objective function of Equation 24 using Equation 25 and the functional form constraint for θ yields:

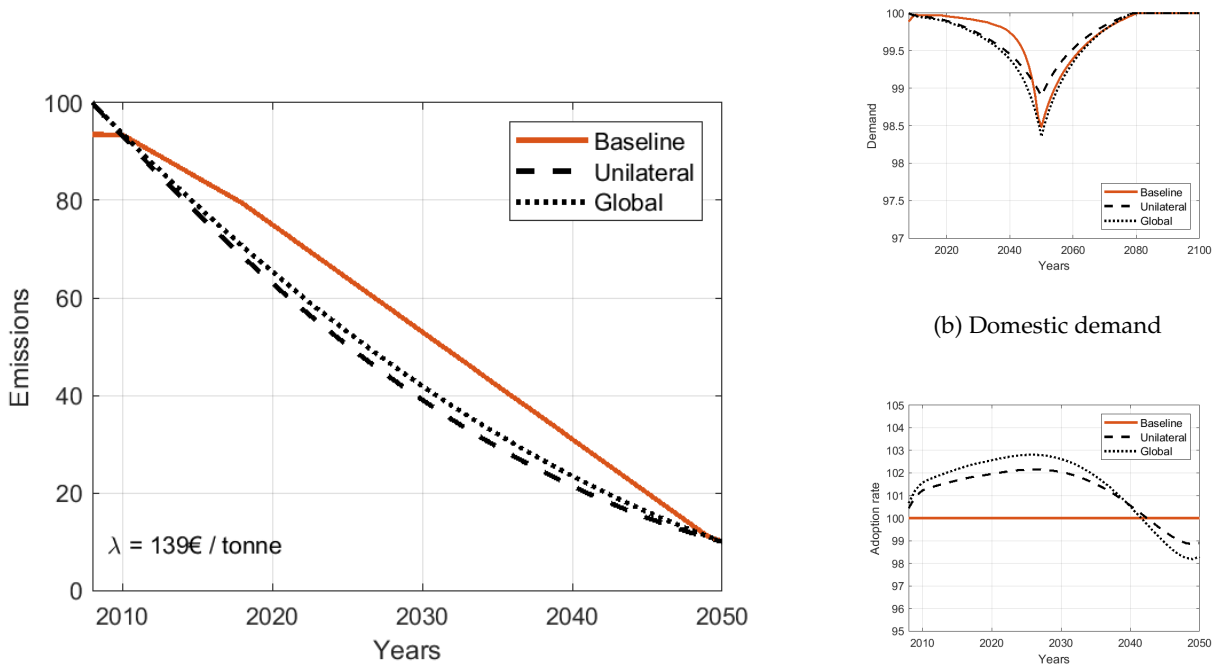
$$\begin{aligned} \max_{b_1} \left\{ \sum_{t=t_0}^{T} \beta_G^t \left[Q_{h,t}(b_1) - \lambda E_t(b_1) \right] + \sum_{t=T+1}^{\infty} \beta_G^t \left[Q_{h,t}(b_1) - \lambda E_t(b_1) \right] \right\} \quad & s.t. \\ E_{h,t} < E_{h,T} \; \forall t < T & \\ E_{h,t} > E_{h,0} \; \forall t < T & \\ E^S_{h,t} = (1 - 0.02) E^S_{h,t-1} \; \forall t > T & \end{aligned} \tag{26}$$

where Equation 26 determines the optimal value of b_1 for which utility by the domestic consumer is maximized given a set of boundary constraints.

6.2 Optimal curvature

To infer the parameter governing the evolution of the emission cap, b_1 in an optimal setting, I consider a social cost of carbon of $\lambda = 156\$/\text{tonne}$, following Farrokhi and Lashkaripour (2022) which translates into 136€/tonne using the exchange rate of 1 January 2019.[22] Unlike firms, the government takes a longer-term perspective on carbon policy, applying a higher discount factor to future utility levels; I set $\beta_G = 0.985$; a common assumption in integrated assessment models (Golosov et al., 2014).

Figure 10: Optimal trajectory of carbon emissions



(a) Optimal path of carbon emissions

(b) Domestic demand

(c) Adoption rate

Note: Evolution of the optimal carbon trajectory across different scenarios, as defined by Equation 26 with $\lambda = 139€/\text{tonne}$. Figure 10a denotes the optimal trajectory of the carbon path for domestic emissions. The dotted line shows the socially optimal path determined by the planner, the red line indicates the current scenario, and the dashed line depicts the optimal path under unilateral action. Figure 10b presents aggregate losses in domestic demand relative to the business-as-usual (BAU) scenario, where BAU = 100%. Finally, Figure 10c depicts the adoption rate relative to the baseline scenario, which is normalized to 100% for each year.

The optimal trajectory of carbon emissions under the unilateral policy scenario is illustrated in Fig-

[22] Subject to additional assumptions, this framework could be extended to integrate into a full Integrated Assessment Model, such as those in Nordhaus (1992) and Golosov et al. (2014) among others, that determines λ endogenously as a function of the stock of emissions. However, this lies beyond the scope of the present paper.

ure 10a. This trajectory represents the path that maximizes the utility of the representative domestic consumer, given the policy constraints imposed by the EU government in Equation 26. In this setting, emissions play a more significant role than in the baseline scenario. By 2050, cumulative emissions in the unilateral case are reduced by 15% relative to the current scenario, while aggregate demand experiences a 0.11% loss in output. This outcome is driven by a higher yet more gradual initial carbon price, which encourages earlier adoption of clean technologies by firms (Figure 10c) which increases by 1%. This in turn mitigates the negative impact of stricter policies on domestic demand in later periods (Figure 10b).

The scenario assumes symmetry in both output and emission output between countries. If the foreign country were to have a higher emission output, the overall emission reduction trajectory would decelerate slightly, as the marginal effectiveness of domestic policies in reducing total emissions would diminish. Further details and results are provided in Appendix G.

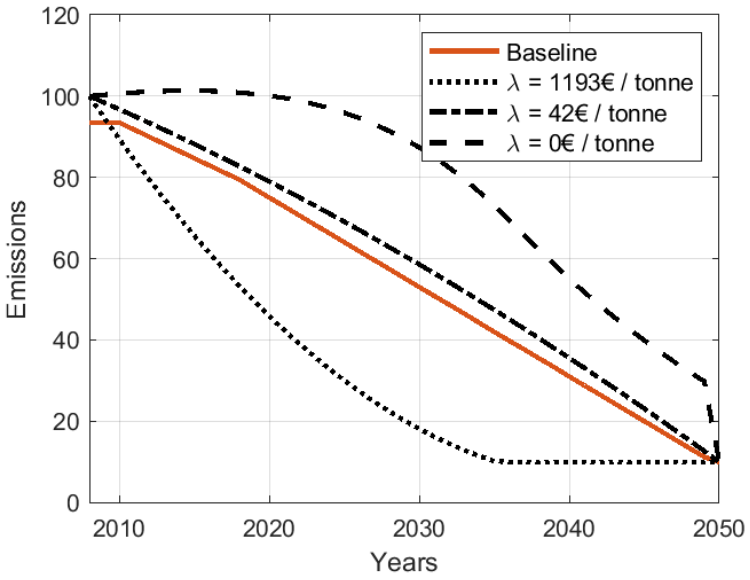
The optimal unilateral policy is evaluated against a first-best benchmark in which a global social planner coordinates emission reductions across both domestic and foreign producers, as well as against the current (status quo) scenario. The carbon trajectory under the coordinated, first-best outcome evolves more gradually than under unilateral action (Figure 10a). This difference arises for three main reasons. First, the marginal effectiveness of policy interventions is greater when both domestic and foreign emissions are reduced simultaneously. Second, under global coordination, domestic consumers can no longer substitute toward cheaper, higher-emission foreign goods. Both imply that a less stringent policy is required to achieve comparable emission targets. The cumulative domestic emissions saved at year 2050 amounts to 12% relative to the baseline scenario, at a cost of 0.18% in terms of aggregate demand. Not having to take into account the asymmetric distortion effects from a trading partner facilitates a more accelerated implementation of the policy. Abatement adoption increases with 0.88%. Since total aggregate emissions decline at the same rate as domestic emissions, even though domestic indicators are worse, aggregate welfare is higher in this situation.

6.3 Alternative costs of carbon

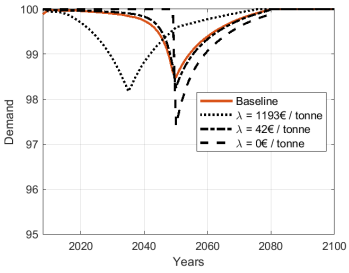
The social cost of carbon used in the previous analysis is used as an average from the literature. Depending on the value of λ, the curvature of the emission path will be different. To see how the curvature changes, I consider two more extreme cases in determining carbon policy. First, I take a value of the social cost of carbon of 1193€/tonne as discussed by Bilal and Kanzig (2024). Second, as a unique case, I consider the situation where the domestic government does not internalize the implications of carbon policy but aims to maximize aggregate demand rather than aggregate welfare ($\lambda = 0$). The government relies on the forward-looking nature of the firm and their profit-maximizing incentives in order to reduce aggregate emissions.[23]

[23] In this scenario, the market clearing condition for the emissions market (Equation 11) is violated as actual emissions are below the allowed emissions and the price of carbon is bounded by $\zeta_t \geq 0$. Market clearing will only be imposed from T onward. Consequently, $\zeta_t = 0$ for $t < T$, prior to achieving net zero. This will in turn generate no output losses prior to net zero and let all emission reductions come through the adoption channel which in turn depends on the firms' expectations. The value of b_1 that maximizes the optimization function of Equation 26 can be obtained analytically: $b_1 \geq 2.15$ which yields $E^S_{h,t} = E_{h,0}$ for $t < T$ and $E^S_{h,T} = E_{h,T}$ at the final year of the

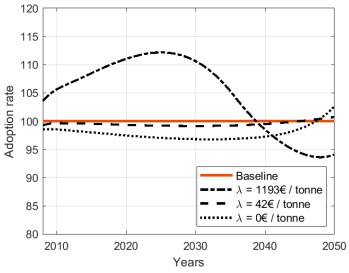
Figure 11: Outcome variables for various trajectories under different λ



(a) Optimal trajectory of carbon emissions



(b) Demand



(c) Adoption rate

Note: Evolution of outcome variables for the optimization of the policy path towards Net zero in 2050 for different values of the social cost of lambda in the welfare functions. The dashed, dot-dash and dotted line line represent $\lambda = \{0, 42, 1193\}$€/tonne. Figure 11a illustrates the optimal trajectory of domestic emissions, representing minimization function (Equation 26). The results are normalized such that the year 2005 equals 100%. Figure 11b presents aggregate losses in domestic demand relative to the business-as-usual (BAU) scenario, where BAU = 100%. Finally, Figure 11c depicts the adoption rate relative to the baseline scenario, which is normalized to 100% for each year.

Figure 11 illustrates the results for the social cost of carbon within the welfare function. As the weight placed on carbon emissions increases, the emission trajectory declines more rapidly toward zero. When the social cost of carbon considered is $\lambda = 1193$ €/tonne, emissions fall sharply with the peak in output losses occurring more rapidly as well. In contrast, the case with $\lambda = 0$ results in no output losses until 2050, with all emission reductions driven solely by the adoption channel. The intermediate case (with $\lambda = 42$ €/tonne) represents a middle ground, closely resembling the baseline scenario.

6.4 Inferring the implied cost of carbon

Finally, the previous section presumes that the government should deviate from their current scenario to an improved one, depending on their evaluation of the cost of emissions. In contrast, this section adopts a different perspective, positing that the government already incorporates this consideration into its max-

constraint. Since time is discrete at the yearly level, b_1 cannot be exactly identified. only the lower bound can be obtained.

imization function. Under the assumption that the government is indeed maximizing domestic welfare, I can back out the implied shadow cost of emissions λ. The value of b_1 that minimizes the differential between $\theta(t)$ and the evolution of the current cap equals $b_1 = -0.0059$.[24] The price of carbon that justifies a value of $b_1 = -0.0059$ equals $\lambda = 81€/\text{tonne}$.

7. Conclusion

This paper highlights the central role of technology adoption in mediating the trade-off between emissions reductions and output losses under unilateral EU climate policy. By combining a dynamic model of heterogeneous firm behavior with detailed panel data, the analysis shows that large, high-emission firms drive most of the observed reductions through capital-intensive abatement investments. This mechanism explains the persistence of low carbon prices in the 2005 - 2020 period. The model quantifies aggregate demand losses, peaking at around 1.9%. By 2025, the adoption channel already mitigated roughly 50% of potential output losses, and by 2050 around 65%.

Further research is needed to relate the evolution in abatement technology with input factors and to see how aggregate welfare is affected through the supply chain. Additionally, it could also shed more light on carbon leakage within multi-plant firms. Finally, the interaction between this model and models of directed technological change such as Acemoglu et al. (2012) could be explored in more detail.

Counterfactual analysis targeting the optimal unilateral trajectory to achieve net zero by 2050 indicates that a moderately faster increase in policy stringency - faster than the current EU trajectory but slower than under full global coordination - can reduce projected domestic cumulative emissions by 15% at a cost of only 0.11% of output in 2050.

[24] This λ does not imply the justification of the carbon policy as a whole, the latter also depends on other constraints. It does measure the pure marginal damage of one extra tonne of CO2 in the welfare function.

References

D. Acemoglu, P. Aghion, L. Bursztyn, and D. Hemous. The Environment and Directed Technical Change. *American Economic Review*, 102(1):131–166, Feb. 2012. ISSN 0002-8282. doi: 10.1257/aer.102.1.131. URL https://www.aeaweb.org/articles?id=10.1257/aer.102.1.131. 5, 31, 36, 99

D. Acemoglu, U. Akcigit, D. Hanley, and W. Kerr. Transition to Clean Technology. *Journal of Political Economy*, 124(1): 52–104, Feb. 2016. ISSN 0022-3808. doi: 10.1086/684511. URL https://www.journals.uchicago.edu/doi/full/10.1086/684511. Publisher: The University of Chicago Press. 21, 105

D. Acemoglu, U. Akcigit, H. Alp, N. Bloom, and W. Kerr. Innovation, Reallocation, and Growth. *American Economic Review*, 108(11):3450–3491, Nov. 2018. ISSN 0002-8282. doi: 10.1257/aer.20130470. URL https://pubs.aeaweb.org/doi/10.1257/aer.20130470. 54

D. A. Ackerberg, K. Caves, and G. Frazer. Identification Properties of Recent Production Function Estimators. *Econometrica*, 83(6):2411–2451, 2015. ISSN 0012-9682. URL https://www.jstor.org/stable/43866416. Publisher: [Wiley, The Econometric Society]. 72

P. Aghion and P. Howitt. A Model of Growth Through Creative Destruction. *Econometrica*, 60(2):323–351, 1992. ISSN 0012-9682. doi: 10.2307/2951599. URL https://www.jstor.org/stable/2951599. Publisher: [Wiley, Econometric Society]. 54

G. Alessandria, C. Arkolakis, and K. J. Ruhl. Firm Dynamics and Trade. *Annual Review of Economics*, 13(Volume 13, 2021):253–280, Aug. 2021a. ISSN 1941-1383, 1941-1391. doi: 10.1146/annurev-economics-090919-025159. URL https://www.annualreviews.org/content/journals/10.1146/annurev-economics-090919-025159. Publisher: Annual Reviews. 4

G. Alessandria, H. Choi, and K. J. Ruhl. Trade adjustment dynamics and the welfare gains from trade. *Journal of International Economics*, 131:103458, July 2021b. ISSN 0022-1996. doi: 10.1016/j.jinteco.2021.103458. URL https://www.sciencedirect.com/science/article/pii/S0022199621000350. 4

P. Antras and D. Chor. *World Trade Evolution: Growth, Productivity and Employment*. Taylor & Francis Group, 2018. Pages: 126-194. 21, 82

C. Arkolakis, A. Costinot, and A. RodrÃguez-Clare. New Trade Models, Same Old Gains? *American Economic Review*, 102(1):94–130, Feb. 2012. ISSN 0002-8282. doi: 10.1257/aer.102.1.94. URL https://www.aeaweb.org/articles?id=10.1257/aer.102.1.94. 21, 50, 81

J. Asker, A. Collard-Wexler, and J. De Loecker. Dynamic Inputs and Resource (Mis)Allocation. *Journal of Political Economy*, 122(5):1013–1063, Oct. 2014. ISSN 0022-3808, 1537-534X. doi: 10.1086/677072. URL https://www.journals.uchicago.edu/doi/10.1086/677072. 18, 19, 71, 72, 73

A. Atkeson and A. Burstein. Innovation, Firm Dynamics, and International Trade. *Journal of Political Economy*, 118(3): 433–484, June 2010. ISSN 0022-3808. doi: 10.1086/653690. URL https://www.journals.uchicago.edu/doi/10.1086/653690. Publisher: The University of Chicago Press. 4, 12

A. Atkeson and P. J. Kehoe. Models of Energy Use: Putty-Putty versus Putty-Clay. *American Economic Review*, 89(4): 1028–1043, Sept. 1999. ISSN 0002-8282. doi: 10.1257/aer.89.4.1028. URL https://www.aeaweb.org/articles?id=10.1257/aer.89.4.1028. 13, 77

G. Barrows and H. Ollivier. Cleaner firms or cleaner products? How product mix shapes emission intensity from manufacturing. *Journal of Environmental Economics and Management*, 88:134–158, Mar. 2018. ISSN 0095-0696. doi: 10.1016/j.jeem.2017.10.008. URL https://www.sciencedirect.com/science/article/pii/S0095069616305083. 5, 11, 57

E. Bartelsman, J. Haltiwanger, and S. Scarpetta. Cross-Country Differences in Productivity: The Role of Allocation and Selection. *American Economic Review*, 103(1):305–334, Feb. 2013. ISSN 0002-8282. doi: 10.1257/aer.103.1.305. URL https://pubs.aeaweb.org/doi/10.1257/aer.103.1.305. 13, 72

S. Basco, G. Felice, B. Merlevede, and M. Mestieri. Financial Crises and the Global Supply Network: Evidence from Multinational Enterprises, May 2023. URL https://www.nber.org/papers/w31216. 46

J. E. Bessen, M. Goos, A. Salomons, and W. Van Den Berge. Automatic Reaction - What Happens to Workers at Firms that Automate? *SSRN Electronic Journal*, 2019. ISSN 1556-5068. doi: 10.2139/ssrn.3328877. URL https://www.ssrn.com/abstract=3328877. 19, 73

A. Bilal and D. R. Kanzig. The Macroeconomic Impact of Climate Change: Global vs. Local Temperature, May 2024. URL https://www.nber.org/papers/w32450. 34, 105

A. Bilal and D. R. Kanzig. Does Unilateral Decarbonization Pay for Itself? *AEA Papers and Proceedings*, 115:369–373, May 2025. ISSN 2574-0768. doi: 10.1257/pandp.20251035. URL https://www.aeaweb.org/articles?id=10.1257/pandp.20251035. 105

A. Bilal and J. H. Stock. A Guide to Macroeconomics and Climate Change, Mar. 2025. URL https://www.nber.org/papers/w33567. 5

N. Bloom, C. I. Jones, J. Van Reenen, and M. Webb. Are Ideas Getting Harder to Find? *American Economic Review*, 110(4):1104–1144, Apr. 2020. ISSN 0002-8282. doi: 10.1257/aer.20180338. URL https://www.aeaweb.org/articles?id=10.1257/aer.20180338. 31

C. Bohringer, K. E. Rosendahl, and H. B. StorrÃžsten. Robust policies to mitigate carbon leakage. *Journal of Public Economics*, 149:35–46, May 2017. ISSN 0047-2727. doi: 10.1016/j.jpubeco.2017.03.006. URL https://www.sciencedirect.com/science/article/pii/S004727271730049X. 4, 28

S. Borghesi, C. Franco, and G. Marin. Outward Foreign Direct Investment Patterns of Italian Firms in the European Union's Emission Trading Scheme*. *The Scandinavian Journal of Economics*, 122(1):219–256, 2020. ISSN 1467-9442. doi: 10.1111/sjoe.12323. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/sjoe.12323. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/sjoe.12323. 5

C. Brunel and A. Levinson. Globalization and Greenhouse Gas Emissions: Evidence from the United States, Jan. 2021. URL https://www.nber.org/papers/w28372. 5

R. J. Caballero, E. M. R. A. Engel, and J. Haltiwanger. Aggregate Employment Dynamics: Building from Microeconomic Evidence. *The American Economic Review*, 87(1):115–137, 1997. ISSN 0002-8282. URL https://www.jstor.org/stable/2950857. Publisher: American Economic Association. 19

B. Callaway, A. Goodman-Bacon, and P. H. C. Sant'Anna. Difference-in-Differences with a Continuous Treatment. *arXiv:2107.02637 [econ]*, July 2021. URL http://arxiv.org/abs/2107.02637. arXiv: 2107.02637. 19, 75

E. Campiglio, S. Dietz, and F. Venmans. Optimal Climate Policy As If the Transition Matters. *SSRN Electronic Journal*, 2022. ISSN 1556-5068. doi: 10.2139/ssrn.4307483. URL https://www.ssrn.com/abstract=4307483. 5

E. Cantillon and A. Slechten. Market Design for the Environment, Dec. 2023. URL https://www.nber.org/papers/w31987. 2, 9, 16, 21

D. Capelle, D. Kirti, N. Pierri, and Villegas-Bauer. Mitigating Climate Change at the Firm Level: Mind the Laggards. *IMF*. URL https://www.imf.org/en/Publications/WP/Issues/2023/11/22/Mitigating-Climate-Change-at-the-Firm-Level-Mind-the-Laggards-541713. 18

J. Cherniwchan and N. Najjar. Do Environmental Regulations Affect the Decision to Export? *American Economic Journal: Economic Policy*, 14(2):125–160, May 2022. ISSN 1945-7731, 1945-774X. doi: 10.1257/pol.20200290. URL https://pubs.aeaweb.org/doi/10.1257/pol.20200290. 4

G. L. Clementi and B. Palazzo. Entry, Exit, Firm Dynamics, and Aggregate Fluctuations. *American Economic Journal: Macroeconomics*, 8(3):1–41, July 2016. ISSN 1945-7707, 1945-7715. doi: 10.1257/mac.20150017. URL https://pubs.aeaweb.org/doi/10.1257/mac.20150017. 19, 74

J. Colmer, R. Martin, M. Muuls, and U. J. Wagner. Does Pricing Carbon Mitigate Climate Change? Firm-Level Evidence from the European Union Emissions Trading Scheme. *SSRN Electronic Journal*, 2020. ISSN 1556-5068. doi: 10.2139/ssrn.3725482. URL https://www.ssrn.com/abstract=3725482. 2, 5, 21

R. W. Cooper and J. C. Haltiwanger. On the Nature of Capital Adjustment Costs. *The Review of Economic Studies*, 73(3):611–633, July 2006. ISSN 0034-6527. doi: 10.1111/j.1467-937X.2006.00389.x. URL https://doi.org/10.1111/j.1467-937X.2006.00389.x. 5, 15, 19, 74

B. R. Copeland and M. S. Taylor. Trade, Growth, and the Environment. *Journal of Economic Literature*, 42(1):7–71, Mar. 2004. ISSN 0022-0515. doi: 10.1257/002205104773558047. URL https://www.aeaweb.org/articles?id=10.1257/002205104773558047. 4, 12, 13

B. R. Copeland, J. S. Shapiro, and M. Scott Taylor. Globalization and the environmentâ. In G. Gopinath, E. Helpman, and K. Rogoff, editors, *Handbook of International Economics*, volume 5 of *Handbook of International Economics: International Trade, Volume 5*, pages 61–146. Elsevier, Jan. 2022. doi: 10.1016/bs.hesint.2022.02.002. URL https://www.sciencedirect.com/science/article/pii/S1573440422000028. 19, 81

J. M. David and V. Venkateswaran. The Sources of Capital Misallocation. *American Economic Review*, 109(7):2531–2567, July 2019. ISSN 0002-8282. doi: 10.1257/aer.20180336. URL https://pubs.aeaweb.org/doi/10.1257/aer.20180336. 19

O. De Jonghe, K. Mulier, and G. Schepens. Going Green by Putting a Price on Pollution: Firm-level Evidence from the EU. *SSRN Electronic Journal*, Working Paper NÂ° 390, 2020. ISSN 1556-5068. doi: 10.2139/ssrn.3725061. URL https://www.ssrn.com/abstract=3725061. 5

A. Dechezlepretre and M. Sato. The Impacts of Environmental Regulations on Competitiveness. *Review of Environmental Economics and Policy*, 11(2):183–206, July 2017. ISSN 1750-6816. doi: 10.1093/reep/rex013. URL https://www.journals.uchicago.edu/doi/10.1093/reep/rex013. Publisher: The University of Chicago Press. 5

A. Dechezlepretre, D. Nachtigall, and F. Venmans. The joint impact of the European Union emissions trading system on carbon emissions and economic performance. 2018. Publisher: OECD. 5

M. A. Drupp, F. Nesje, and R. C. Schmidt. Pricing Carbon: Evidence from Expert Recommendations. *American Economic Journal: Economic Policy*, 16(4):68–99, Nov. 2024. ISSN 1945-7731. doi: 10.1257/pol.20220571. URL https://www.aeaweb.org/articles?id=10.1257/pol.20220571. 1

A. Dube, D. Girardi, Ã. JordÃ , and A. M. Taylor. A Local Projections Approach to Difference-in-Differences, Apr. 2023. URL https://www.nber.org/papers/w31184. 52

F. Farrokhi and A. Lashkaripour. Can Trade Policy Mitigate Climate Change? 2022. 4, 13, 28, 32, 33

G. Felbermayr, S. Peterson, and J. Wanner. Trade and the environment, trade policies and environmental policiesâHow do they interact? *Journal of Economic Surveys*, 39(3):1148–1184, 2025. ISSN 1467-6419. doi: 10.1111/joes.12628. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/joes.12628. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/joes.12628. 5

A. Finkelstein Shapiro and G. E. Metcalf. The macroeconomic effects of a carbon tax to meet the U.S. Paris agreement target: The role of firm creation and technology adoption. *Journal of Public Economics*, 218:104800, Feb. 2023. ISSN 0047-2727. doi: 10.1016/j.jpubeco.2022.104800. URL https://www.sciencedirect.com/science/article/pii/S004727272200202X. 18

L. Fontagne and K. Schubert. The Economics of Border Carbon Adjustment: Rationale and Impacts of Compensating for Carbon at the Border. *Annual Review of Economics*, 15(1):389–424, 2023. doi: 10.1146/annurev-economics-082322-034040. URL https://doi.org/10.1146/annurev-economics-082322-034040. _eprint: https://doi.org/10.1146/annurev-economics-082322-034040. 4, 28

R. Forslid, T. Okubo, and K. H. Ulltveit-Moe. Why are firms that export cleaner? International trade, abatement and environmental emissions. *Journal of Environmental Economics and Management*, 91:166–183, Sept. 2018. ISSN 0095-0696. doi: 10.1016/j.jeem.2018.07.006. URL https://www.sciencedirect.com/science/article/pii/S0095069617303340. 4

L. Foster, J. Haltiwanger, and C. Syverson. Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability? *American Economic Review*, 98(1):394–425, Feb. 2008. ISSN 0002-8282. doi: 10.1257/aer.98.1.394. URL https://pubs.aeaweb.org/doi/10.1257/aer.98.1.394. 57, 71, 73

M. Fowlie, M. Reguant, and S. P. Ryan. Market-Based Emissions Regulation and Industry Dynamics. *Journal of Political Economy*, 124(1):249–302, Feb. 2016. ISSN 0022-3808. doi: 10.1086/684484. URL https://www.journals.uchicago.edu/doi/full/10.1086/684484. Publisher: The University of Chicago Press. 5

M. L. Fowlie and M. Reguant. Mitigating Emissions Leakage in Incomplete Carbon Markets. *Journal of the Association of Environmental and Resource Economists*, 9(2):307–343, Mar. 2022. ISSN 2333-5955. doi: 10.1086/716765. URL https://www.journals.uchicago.edu/doi/full/10.1086/716765. Publisher: The University of Chicago Press. 4, 28

M. Golosov, J. Hassler, P. Krusell, and A. Tsyvinski. Optimal Taxes on Fossil Fuel in General Equilibrium. *Econometrica*, 82(1):41–88, 2014. ISSN 1468-0262. doi: 10.3982/ECTA10217. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA10217. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA10217. 5, 12, 13, 14, 33

L. H. Goulder and A. Schein. Carbon Taxes vs. Cap and Trade: A Critical Review, Aug. 2013. URL https://papers.ssrn.com/abstract=2311587. 6

J. Haltiwanger, R. S. Jarmin, and J. Miranda. WHO CREATES JOBS? SMALL VERSUS LARGE VERSUS YOUNG. *The Review of Economics and Statistics*, 95(2):347–361, 2013. ISSN 0034-6535. URL https://www.jstor.org/stable/43554390. Publisher: The MIT Press. 54

J. Haltiwanger, R. S. Jarmin, R. Kulick, and J. Miranda. High Growth Young Firms: Contribution to Job, Output, and Productivity Growth. In *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges*, pages 11–62. University of Chicago Press, Sept. 2016. URL https://www.nber.org/books-and-chapters/measuring-entrepreneurial-businesses-current-knowledge-and-challenges/high-growth-young-firms-contribution-job-output-and-productivity-growth. 54

J. Hassler, P. Krusell, and C. Olovsson. Directed Technical Change as a Response to Natural Resource Scarcity. *Journal of Political Economy*, 129(11):3039–3072, Nov. 2021. ISSN 0022-3808. doi: 10.1086/715849. URL https://www.journals.uchicago.edu/doi/full/10.1086/715849. Publisher: The University of Chicago Press. 5

K. Head and J. Ries. Increasing Returns Versus National Product Differentiation as an Explanation for the Pattern of U.S.âCanada Trade. *American Economic Review*, 91(4):858–876, Sept. 2001. ISSN 0002-8282. doi: 10.1257/aer.91.4.858. URL https://pubs.aeaweb.org/doi/10.1257/aer.91.4.858. 21, 71, 82

D. Hemous. The dynamic impact of unilateral environmental policies. *Journal of International Economics*, 103:80–95, Nov. 2016. ISSN 0022-1996. doi: 10.1016/j.jinteco.2016.09.001. URL https://www.sciencedirect.com/science/article/pii/S0022199616301052. 5, 31

C.-T. Hsieh and P. J. Klenow. Misallocation and Manufacturing TFP in China and India. *The Quarterly Journal of Economics*, 124(4):1403–1448, 2009. ISSN 0033-5533. URL https://www.jstor.org/stable/40506263. Publisher: Oxford University Press. 13, 14

J. JaraitÄ, T. Jong, A. KaÅŸukauskas, A. Zaklan, and A. Zeitlberger. Ownership links and enhanced EUTL dataset, 2016. URL https://cadmus.eui.eu/handle/1814/64596. Accepted: 2019-10-16T15:43:28Z Publisher: European University Institute, RSCAS, FSR. 44

C. I. Jones. The Past and Future of Economic Growth: A Semi-Endogenous Perspective. *Annual Review of Economics*, 14(1):null, 2022. doi: 10.1146/annurev-economics-080521-012458. URL https://doi.org/10.1146/annurev-economics-080521-012458. _eprint: https://doi.org/10.1146/annurev-economics-080521-012458. 31

S. Kalemli-Ozcan, B. Sorensen, C. Villegas-Sanchez, V. Volosovych, and S. Yesiltas. How to Construct Nationally Representative Firm Level Data from the Orbis Global Database: New Facts and Aggregate Implications. Technical Report w21558, National Bureau of Economic Research, Cambridge, MA, Sept. 2015. URL http://www.nber.org/papers/w21558.pdf. 10, 45, 47

T. Klette and S. Kortum. Innovating Firms and Aggregate Innovation. *Journal of Political Economy*, 112(5):986–1018, Oct. 2004. ISSN 0022-3808. doi: 10.1086/422563. URL https://www.journals.uchicago.edu/doi/abs/10.1086/422563. Publisher: The University of Chicago Press. 54

S. Kortum and D. Weisbach. Optimal Unilateral Carbon Policy. *Coase-Sandor Working Paper Series in Law and Economics*, Jan. 2021. URL https://chicagounbound.uchicago.edu/law_and_economics/957. 4, 28, 32

A. Lanteri and A. Rampini. Financing the Adoption of Clean Technology. 2025. 18

S. Letout. Corporate micro-data in the EU Emission Trading System (EU ETS). *Publications Office of the European Union*, EUR 31002 EN, Mar. 2022. ISSN 1831-9424. doi: 10.2760/77710. URL https://publications.jrc.ec.europa.eu/repository/handle/JRC127805. 10, 45

R. Martin, M. Muuls, L. B. de Preux, and U. J. Wagner. Industry Compensation under Relocation Risk: A Firm-Level Analysis of the EU Emissions Trading Scheme. *American Economic Review*, 104(8):2482–2508, Aug. 2014. ISSN 0002-8282. doi: 10.1257/aer.104.8.2482. URL https://pubs.aeaweb.org/doi/10.1257/aer.104.8.2482. 9, 14

R. Martin, M. Muuls, and U. J. Wagner. 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*, 10(1):129–148, Jan. 2016. ISSN 1750-6816. doi: 10.1093/reep/rev016. URL https://www.journals.uchicago.edu/doi/10.1093/reep/rev016. Publisher: The University of Chicago Press. 5

M. J. Melitz. The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica*, 71(6):1695–1725, 2003. ISSN 0012-9682. URL https://www.jstor.org/stable/1555536. Publisher: [Wiley, Econometric Society]. 4, 12

M. J. Melitz and S. Polanec. Dynamic Olley-Pakes productivity decomposition with entry and exit. *The RAND Journal of Economics*, 46(2):362–375, June 2015. ISSN 07416261. doi: 10.1111/1756-2171.12088. URL https://onlinelibrary.wiley.com/doi/10.1111/1756-2171.12088. 11, 50, 56, 57

G. E. Metcalf. Five Myths About Carbon Pricing, Apr. 2023. URL https://www.nber.org/papers/w31104. 6

N. Najjar and J. Cherniwchan. Environmental Regulations and the Cleanup of Manufacturing: Plant-Level Evidence. *The Review of Economics and Statistics*, 103(3):476–491, July 2021. ISSN 0034-6535. doi: 10.1162/rest_a_00904. URL https://doi.org/10.1162/rest_a_00904. 5, 11, 57

W. D. Nordhaus. An Optimal Transition Path for Controlling Greenhouse Gases. *Science*, 258(5086):1315–1319, 1992. ISSN 0036-8075. URL https://www.jstor.org/stable/2880417. Publisher: American Association for the Advancement of Science. 5, 33, 105

M. Pahle, S. Quemin, S. Osorio, C. GÃŒnther, and R. Pietzcker. The emerging endgame: The EU ETS on the road towards climate neutrality. *Resource and Energy Economics*, 81:101476, Feb. 2025. ISSN 0928-7655. doi: 10.1016/j.reseneeco.2024.101476. URL https://www.sciencedirect.com/science/article/pii/S0928765524000526. 9, 17, 24

N. Ramondo, A. RodrÃguez-Clare, and F. Tintelnot. Multinational Production: Data and Stylized Facts. *American Economic Review*, 105(5):530–536, May 2015. ISSN 0002-8282. doi: 10.1257/aer.p20151046. URL https://www.aeaweb.org/articles?id=10.1257/aer.p20151046. 46, 102

D. Restuccia and R. Rogerson. Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics*, 11(4):707–720, Oct. 2008. ISSN 1094-2025. doi: 10.1016/j.red.2008.05.002. URL https://www.sciencedirect.com/science/article/pii/S1094202508000203. 13, 14

E. Rottner and K. von Graevenitz. What Drives Carbon Emissions in German Manufacturing: Scale, Technique or Composition? *Environmental and Resource Economics*, 87(9):2521–2542, Sept. 2024. ISSN 1573-1502. doi: 10.1007/s10640-024-00894-7. URL https://doi.org/10.1007/s10640-024-00894-7. 5, 11, 57

S. P. Ryan. The Costs of Environmental Regulation in a Concentrated Industry. *Econometrica*, 80(3):1019–1061, 2012. ISSN 1468-0262. doi: 10.3982/ECTA6750. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA6750. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA6750. 5

J. S. Shapiro. The Environmental Bias of Trade Policy*. *The Quarterly Journal of Economics*, 136(2):831–886, May 2021. ISSN 0033-5533. doi: 10.1093/qje/qjaa042. URL https://doi.org/10.1093/qje/qjaa042. 4, 11, 28, 32

J. S. Shapiro and R. Walker. Why Is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade. *American Economic Review*, 108(12):3814–3854, Dec. 2018. ISSN 0002-8282. doi: 10.1257/aer.20151272. URL https://www.aeaweb.org/articles?id=10.1257/aer.20151272. 4, 5, 12, 13, 19, 57, 81

V. Sterk, P. SedlÃ¡Äek, and B. Pugsley. The Nature of Firm Growth. *American Economic Review*, 111(2):547–579, Feb. 2021. ISSN 0002-8282. doi: 10.1257/aer.20190748. URL https://pubs.aeaweb.org/doi/10.1257/aer.20190748. 13, 46

L. Sun and S. Abraham. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199, Dec. 2021. ISSN 03044076. doi: 10.1016/j.jeconom.2020.09.006. URL https://linkinghub.elsevier.com/retrieve/pii/S030440762030378X. 19

G. R. Timilsina. Carbon Taxes. *Journal of Economic Literature*, 60(4):1456–1502, Dec. 2022. ISSN 0022-0515. doi: 10.1257/jel.20211560. URL https://www.aeaweb.org/articles?id=10.1257/jel.20211560. 29

I. Van Beveren. Total Factor Productivity Estimation: A Practical Review. *Journal of Economic Surveys*, 26(1):98–128, 2012. ISSN 1467-6419. doi: 10.1111/j.1467-6419.2010.00631.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-6419.2010.00631.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-6419.2010.00631.x. 72

F. van der Ploeg and A. Rezai. The risk of policy tipping and stranded carbon assets. *Journal of Environmental Economics and Management*, 100:102258, Mar. 2020. ISSN 0095-0696. doi: 10.1016/j.jeem.2019.102258. URL https://www.sciencedirect.com/science/article/pii/S0095069618302481. 5

S. F. Verde. The Impact of the Eu Emissions Trading System on Competitiveness and Carbon Leakage: The Econometric Evidence. *Journal of Economic Surveys*, 34(2):320–343, 2020. ISSN 1467-6419. doi: 10.1111/joes.12356. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/joes.12356. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/joes.12356. 5

A. Young. Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results*. *The Quarterly Journal of Economics*, 134(2):557–598, May 2019. ISSN 0033-5533. doi: 10.1093/qje/qjy029. URL https://doi.org/10.1093/qje/qjy029. 19, 55, 75

A. Data gathering and analysis

A.1 Data gathering process

This section describes the construction of the dataset and emphasizes the economic relevance of the firms included within the European Union. The primary dataset is assembled by combining two main sources: emissions data for all firms participating in the EU Emissions Trading System (ETS), obtained from the European Union Transaction Log (EUTL), and annual financial information from ORBIS, provided by Bureau van Dijk.

The section first outlines the collection of emissions data, followed by the process of obtaining the financial information. It then details the methodology used to merge these two datasets and to construct several key variables. Finally, it introduces additional data sources employed in the study and explains how they are incorporated into the model used in the analysis.

A.1.1 Emission data

The data collection process begins with extracting information from the European Union Transaction Log (EUTL) for the period 2005â2022. The preparation of these data follows the methodology described in JaraitÄ et al. (2016). The EUTL records both the allocated and surrendered emission allowances at the installation level for all firms operating within the EU. Moreover, it provides annual ownership details, including the address and company registration number of each installation, which are essential for linking it to the financial statements of the firm.

Table A.1 shows the processes to obtain the baseline dataset used in the paper. The source data encompasses all information on all installations found on the EUTL (19624 installations). Airline companies and entities having an account, but lacking an installation identifier in the EUTL are excluded. Such latter accounts are typically utilized solely for the purpose of trading carbon certificates. All installations are aggregated up to the firm level. I define 'a firm' as an entity with a unique company registration number in the database.

In cases where a firm possesses multiple accounts but at least one of these accounts is linked to an installation that omits emissions, the emission certificates used will be aggregated up to the firm level. Installations associated with more than one firm in a given year are assigned to the firm that retains ownership in the subsequent year. Such cases represent less than 0.1% of all installations. The final dataset comprises 10,676 firms.

A first indicator is developed at the firm level that tracks the evolution of entry and exit of installations over time. Note that entry of the installation (and therefore potentially also at the firm level) in the dataset might be involuntary and depends on the stringency of the regulations set forward by the European Union. In 2013, a policy revision expanded the range of sectors and firms covered by the EU ETS, leading to a substantial increase in the number of new installations. This expansion could introduce bias into the aggregate results if not properly addressed. The development of this indicator enables seeing when the installation was first observed in the data but also when it was acquired by a specific firm. Installations that were not

in the EUTL prior to 2013 are omitted from the analysis. Firms that saw an increase in the number of their installations exposed to the policy change which were not acquired are flagged and omitted in some parts of the analysis (Table A.1)

A.1.2 Financial data

The emission data is combined with the yearly non-consolidated financial accounts of each firm found in ORBIS for the year 2005 to 2020. The data includes, among others, the added value, number of employees, materials, and (fixed) tangible assets. The database also provides information on the date of incorporation, the legal status, as well as the main sector the firm operates in. The analysis is limited to the manufacturing sector due to the better documentation of the financial variables in ORBIS; the relative ease of comparability between firms; and the increased likelihood of carbon leakage relative to for instance firms operating in energy-generating industries. Firms are considered to be manufacturing firms if they have at least one installation that is active in the manufacturing sector (Nace rev.2 2d >9 & ≤ 33)

The matching procedure linking the emissions data (EUTL) and financial data (ORBIS) in this study closely follows the methodology of Letout (2022), with minor manual adjustments. Specifically, in cases where the original automated procedure fails, firms are matched based on clear correspondences in installation names and addresses, verified with the online ORBIS database. Such manual matching is primarily required for German firms (see Table A.1).

Table A.1: Data processing ORBIS global

Method	id	Count	Percentage (%)
Source data	Installations	19624	100%
Omittance empty company registration numbers and airlines	Installations	17176	87.5%
Collapse to company registration number	Firm	10676	100%
Matching procedure Letout (2022)	Firm	8615	80.6%
Manual corrections	Firm	8863	83.0%
Limit to manufacturing sector	Firm	5187	100%
Financial data available	Firm	4220	87.1%
No acquisition operating installation	Firm	3107	87.1%

Note: This table outlines the steps involved in constructing the database used for the analysis. The Method column briefly summarizes the main processing approach. Id, Count, and Percentage refer to the type of unique identifier used and the number and proportion of observations retained at each step. The dataset begins with 19,624 installations from the EUTL, which are consolidated into 10,682 firms. These are then matched with ORBIS data, resulting in 8,863 matched firms. Finally, the analysis is restricted to the manufacturing sector, yielding 4,220 firms with sufficient financial information.

To ensure the financial data is suitable for analysis, I follow the procedures outlined in Kalemli-Ozcan et al. (2015). However, the concerns raised in that study are less critical in the context of this paper, as the

focus is primarily on large firms, which are typically well represented in national accounts, and there is intention in this paper to conduct cross-country comparisons. Except for one table (Table A.4), this study does not analyze cross-country variation, and thus differences in data gathering practices across European countries are not a central concern. As shown in Table A.1, financial information has been successfully identified for nearly 88% of the initial 5,186 manufacturing firms. Firms located in Iceland, the Netherlands, Estonia, Croatia, Greece, Lithuania, and Malta are excluded from the analysis due to insufficient financial reporting. Finally, firms from Norway have no reported emissions for the period from 2005 to 2007.

A.1.3 Deflating financial values

Financial values are deflated using annual sector-country level price indices from Eurostat. Different deflators are applied depending on the variable: material costs and operating revenue are deflated using the producer price index; tangible capital is deflated using a capital goods price index; value added is deflated using the gross value added index. All deflators are normalized to 2015 = 100.

A.1.4 Ownership structure

The ownership linkages are needed to construct an indicator of the multinational character of the firm (Ramondo et al., 2015; Basco et al., 2023). To evaluate the multinational nature of the firms in the dataset, I rely on the ownership linkages at the year 2019, as reported in ORBIS. A firm is considered to own another if it possesses more than 50% of the ownership shares. I assume that ownership linkages remain constant over time.

I distinguish between firms that are subsidiaries of other entities and those that are owners of other entities. Additionally, I break down these ownership linkages further, into categories based on their scope, identifying whether they are national, European, or foreign ownership (i.e. non-European) linkages, depending on the nationality of the global ultimate owner of the firm. Shares of the linkages for each category can be found in **??**.

A.1.5 Documentation of data

The main summary statistics in the data are presented in the main body of the paper. Here, I only add a comparison table that accounts for the exclusion restriction of the firms that had novel installations after 2012.

The firms that are present in the EU ETS only represent a small subset of all firms active in the total economy of the European Union - there are only 4220 firms in this dataset. However, they account for ≈40% of all emissions and 8% of all GDP, while . The ETS captures mainly large firms, the average firm has more than 1039 employees and adds a yearly value to the economy of more than 633 million euros. The economy is not dynamic, with a yearly entry and exit rate of around 1 percent which is significantly lower than the average country-level (Sterk et al., 2021).

Table A.2: Comparison of firms omitted after 2012 and total

Variable	*Non-omitted*		*Total*	
Number of firms	3106		4220	
Number of employees	905	(5276)	1039	(5445)
Emissions (tonne)	174	(738)	203	(848)
Value added (million Euro)	98.3	(712)	114.0	(744)

Note: This figure compares firms participating in the EU Emissions Trading System (ETS) that are classified as carbon-leakage-prone by the European Commission with those that are not. The number of firms reflects the total count, while the figures for employees (number), emissions (tonne), and value added (in million Euros) represent average values per firm within each group.

A.1.6 Aggregate summary statistics

To provide an overview of the firms included in the analysis, this section presents descriptive statistics of the main dataset. The tables below summarize key characteristics at both the sectoral and country levels, including emissions, value added, and their emission intensity. These statistics offer a first indication of the economic and environmental heterogeneity across firms and regions within the EU ETS sample.

It is important to note that direct cross-country comparisons should be interpreted with caution, as the measurement of certain variables may differ slightly across countries (Kalemli-Ozcan et al., 2015). The resulting figures should therefore be viewed primarily as indicative. Tables Table A.3 and Table A.4 further illustrate how the aggregate statistics are defined and specify the countries and sectors included in the dataset.

A.2 Other data sources

This section outlines additional data sources that supplement the core empirical analysis. These datasets help capture additional empirical robustness tests or are used to parameterize the structural model or any extensions.

A.2.1 price data

The study incorporates annual historical data on carbon prices from the EU Emission Allowances Index and crude oil prices from the Federal Reserve Bank of St. Louis. These timeseries data are used to test whether carbon policy, identified through the price of carbon certificates, has had a marginal effect on the decisions to adopt novel clean technologies rather than other aggregate forces such as the price of crude oil.

Table A.3: Summary statistics between sectors

Nace 2 rev.2	Number of firms	Added Value	Emissions	Emission Intensity
15	4	.02	0	.95
14	5	.05	.01	1.01
31	6	.07	0	.3
33	9	.15	.07	2.23
18	7	.16	.03	.78
32	13	.33	.44	6.68
13	87	.55	.17	1.59
27	27	.71	.57	3.99
26	14	.79	.04	.25
16	90	.9	.88	4.93
25	41	1.6	2.38	7.5
12	10	1.62	.03	.08
22	58	1.98	.41	1.05
11	75	3.31	.22	.33
24	145	4.75	15.95	16.88
30	34	5.45	.13	.12
17	485	5.56	5.36	4.84
10	354	6.1	3.56	2.94
19	73	6.36	22.76	17.98
20	268	7.45	6.1	4.11
21	54	7.59	.34	.23
28	43	7.84	.11	.07
23	944	9.44	39.39	20.98
29	55	27.23	1.05	.19

Note: The table reports average statistics over the period 2005â2019. CO2 emissions are measured in tonnes, and value added is expressed in million euros. Emission intensity is calculated as the ratio of emissions to value added, multiplied by 10,000. The table is ordered by value added, and the number of firms represents the total count of ETS participants within each sector. Sectors are identified according to their NACE Rev. 2 2-digit codes.

A.2.2 Carbon leakage sectors

Second, to identify sectors vulnerable to carbon leakage, the EU established a Carbon Leakage List. This list designates industries at significant risk of relocating to countries with weaker climate policies, due to high carbon costs and strong exposure to international competition. Introduced in 2009, it is updated for each phase of the EU ETS. Sectors qualify through quantitative criteriaâtrade intensity above 17% and/or carbon cost intensity exceeding 0.2% of gross value addedâor through a qualitative assessment demon-

Table A.4: Summary statistics between countries

Country	Number of firms	Added Value	Emissions	Emission Intensity
LV	15	0	.02	16.6
LU	5	.01	.11	23.51
BG	64	.28	.85	5.12
NL	55	.37	1.32	5.89
SK	66	.43	1.77	6.97
IE	12	.47	.55	1.98
SI	61	.57	.52	1.53
PL	194	1.12	4.83	7.21
RO	92	1.21	1.98	2.74
PT	155	1.21	2.96	4.1
FI	46	1.3	1.58	2.05
CZ	139	1.56	5.04	5.42
HU	74	2.03	2.03	1.68
SE	93	2.11	1.73	1.37
AT	70	2.58	4.96	3.22
IT	313	6.69	12.3	3.08
BE	162	8.64	5.18	1.01
GB	191	9.47	4.73	.84
ES	509	9.63	15.42	2.68
FR	302	12.07	16.62	2.31
DE	283	38.26	15.51	.68

Note: average statistics over the period from 2005 to 2019. The CO2 emissions are documented per tone and Value added data in million euros. The multinational dimension is split in whether the firm is a subsidiary of either a domestic, European, or non-European entity - or whether they are the owner of such entities. Note that this sums up to less than 1 as firms without any subsidiary entities can also exist. Entry and Exit rate are computed by dividing all observations relative to the number of observed entry and exits respectively.

strating high emissions intensity and limited cost pass-through ability. Examples include steel, cement, aluminum, and certain chemical industries, which may receive up to 100% of their emission allowances for free. I use this list to generate a 'carbon-leakage'-prone sector in the model. For space constraints, the total list is omitted from the paper. One can find it on https://climate.ec.europa.eu/eu-action/carbon-markets/eu-emissions-trading-system-eu-ets/free-allocation/carbon-leakage_en

Table A.5: Comparison of ETS Firms by Carbon Leakage Status

Variable	*Leakage*	*Non-Leakage*
Number of firms	704	2197
Number of employees	498	1061
Emissions (tonne)	246.4	162.5
Value added (million Euro)	61.7	120.2

Note: comparison between firms participating in the EU Emissions Trading System (ETS) that are classified as carbon-leakage-prone by the European Commission and those that are not. The number of firms reflects the total count, while the figures for employees (number), emissions (tonne), and value added (in million Euros) represent average values per firm within each group.

A.2.3 Aggregate trade data

The iceberg trade-cost from the European Union to the rest of the world is estimated using aggregate data from the World Input Output Tables (WIOT) for 2005 to 2007, following Arkolakis et al. (2012).

B. Reduced-form Evidence

The economic theory presented in the paper suggests that more productive, polluting firms are more likely to adopt novel technologies. The latter because they stand to gain the most from adopting new technology, and the prior because they have the economies of scale to recoup the sunk cost of adoption. To examine (i) whether this pattern appears in the data and (ii) if it serves as a primary driver of the aggregate evolution of emissions in the data, this section provides a series of reduced-form tests.

I start by analyzing the aggregate implications of this theory - as more productive, polluting firms adopt more technologies, inequality in the clean technology margin is decreasing while inequality in size is either stagnant or even increasing. However, this result still holds room for green creative destruction as a potential channel. To ensure that the aggregate results are driven by polluting incumbent firms, I apply a local-projection analysis for the start of each of the two ETS-phases in the data (i.e. phase II in 2008 and phase III in 2013). Finally, I rule out potential relocation to either completely new firms or to firms not covered by the ETS. Finally, I apply a dynamic Olley-Pakes decomposition (Melitz and Polanec, 2015) to ensure the importance of either component to the increase in aggregate clean technology.

B.1 Aggregate statistics

Economic theory states that as large, productive firms adopt more technologies, the disparity on the clean technology margin across firms is narrowing over time whereas the evolution of value added is either increasing or, at best, stagnant. To ensure the consistency of economic theory with the aggregate data and

Figure B.1: Relative evolution inequality



(a) HHI

(b) Market share (C10)

(c) Skewness

(d) Standard deviation

Note: Aggregate trends in various metrics of cross-sectional inequality from 2007 to 2019, the final year before the onset of Covid-19. The red line represents the trajectory of aggregate clean technology, while the black line represents value added. Dashed lines indicate the respective linear trends.

rule out spurious findings. I examine various indicators of inequality, specifically the standard deviation skewness of the distribution, and the Herfindahl-Hirschman Index (HHI) and market share of the 10 largest firms as two nonparametric measures of market concentration.

The aggregate figures presented in the main text employ the standard deviation as a metric of cross-sectional inequality to illustrate that the variation in emissions intensity is declining over time, whereas the evolution of value added is either increasing or, at best, remaining stable. While the standard deviation is a common measure of inequality, it represents only one dimension. To assess the robustness of this relationship and rule out spurious findings, I also examine alternative indicators of inequality, specifically the skewness of the distribution, and the Herfindahl-Hirschman Index (HHI) and market share of the 10

largest firms as two nonparametric measures of market concentration.

The aggregate results are found in Figure B.1. The upper left panel (Figure B.4a) displays the evolution of the Herfindahl-Hirschman Index (HHI) for emissions and value added between 2007 and 2019. The HHI for emissions shows a steady decline over time, indicating a decreasing concentration of emissions among firms, and thus a more even distribution of emissions across the firm population. In contrast, the HHI for value added remains relatively stable, suggesting that market concentration in terms of economic output has not substantially changed over the same period. The same relation holds for the market share of the 10 largest firms (Figure B.4b).

The lower two panels consider the parametric distribution of the logarithm for both variables. Figure B.1c illustrates the skewness of emissions intensity and value added over time. Skewness in emissions intensity is increasing slightly, highlighting that the the share of extreme polluters is declining and that the distribution normalizes. Conversely, the skewness of value added declines, particularly in the early years, indicating an increase in asymmetry towards one where large firms have become more important. Finally, Figure B.1d indicates a converging standard deviation for abatement technology, highlighting a *catch-up mechanism* in the data. This convergence is not present in value added; in contrast, it is even diverging up to a certain point.

In conclusion, these figures indicate that between 2007 and 2019, emissions became more evenly distributed across firms (decreasing HHI and slightly increasing skewness), while market concentration in value added remained stable (stable HHI) with a slight increase in the dominance of large firms (decreasing skewness).

B.2 Local projections

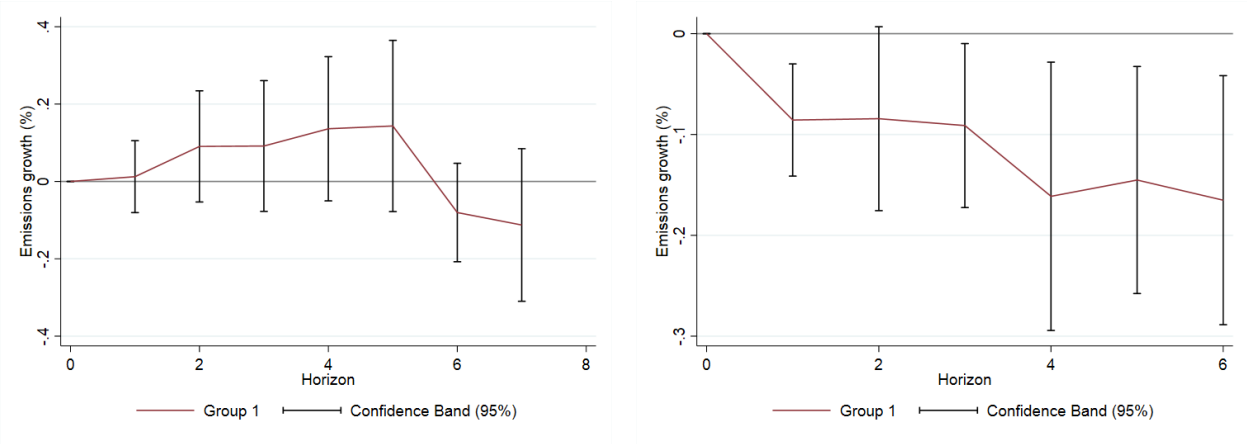
To test whether these results are truly indicating the growth of these incumbents and do not hide a potential shift-share mechanism where small green firms increase their market share, I examine the evolution of incumbents based on their ex-ante conditions. I generate a binary variable $D_{i,t}$ that indicates whether a firm is polluting more than the average firm and larger than the average firm (i.e. $D_{i,t} = 1\{\frac{e_{i,t}}{va_{i,t}} > \frac{\bar{e_{i,t}}}{va_{i,t}}, va_{i,t} > \bar{va_{i,t}}\}$. I then track the evolutions of emissions and size, proxied by value added, over time using a local-projections-type regression (Dube et al., 2023) based on the ex-ante characteristics of the firms prior to the introduction of a novel phase of the EU ETS. The regression looks as follows:

$$\Delta \ln(e_{i,t+h}) = \beta D_{i,t} + \xi_c + \iota_s + \varepsilon_{i,t} \tag{B.1}$$

where $\Delta \ln(e_{i,t+h})$ is the differential in emissions. ξ_c and ι_s are country and sector fixed effects and the variables are indexed with $t \in \{2007, 2012\}$, which either captures the initial conditions at the start of phase II or phase III of the ETS. The results are shown in Figure B.2.

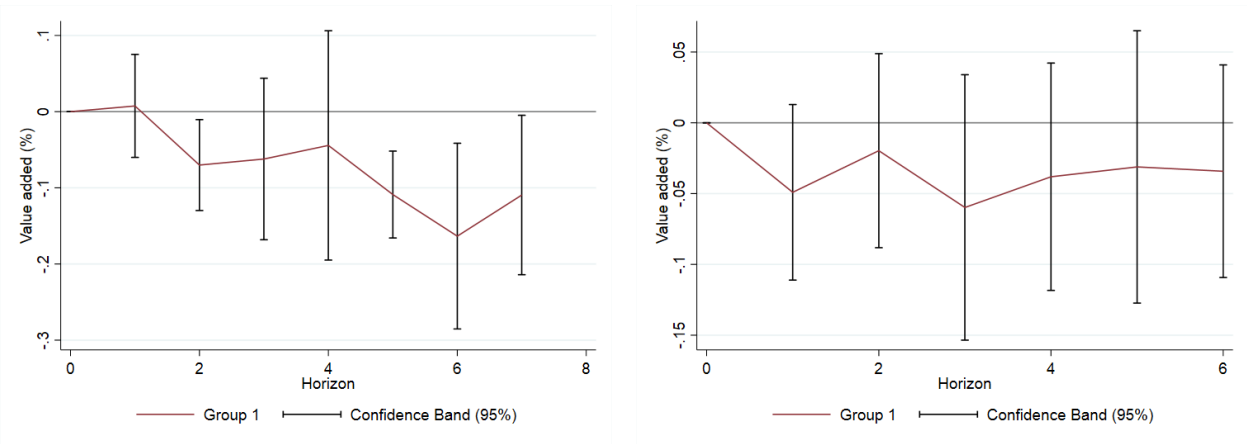
Figure B.2a and Figure B.2c show no significant divergence in emissions or value added between large polluting firms and other firms from the start of 2007. Their evolution coincides with the start of financial crisis and at the start of Phase II of the ETS, It likely reflects the broader economic slowdown, which seems to hint towards large emission-intensive firms being hit more than the average firm. Adding towards this,

Figure B.2: Relative evolution large polluting firms



(a) Emissions (2007)

(b) Emissions (2012)



(c) Value added (2007)

(d) Value added (2012)

Note: Comparative analyses of the growth path of large emission-intensive firms relative to all other firms in the ETS in their evolution of emission output (upper 2 figures) and value added (lower 2 figures) for the year 2007 and 2012. The years are selected as the 2 years prior to the initialisation of a new phase of the ETS. Note that the start of phase III is considered to have been much more stringent that phase II.

the price of carbon has been structurally low throughout this period.

In contrast, from 2012 onward, when aggregate economic conditions had been relatively stabilized, and ahead of the stricter Phase III, emissions among large polluters began to decline more visibly. Figure B.2b shows a downward trend in emissions growth, especially in the first few years after 2012. The point estimate at Horizon 0 is about -10%, and the decline continues gradually to about -15% by 4 years. Note that these point-estimates have large confidence intervals.

Finally, the Figure B.2d shows stable value added, implying emissions reductions occurred without

harming firm performance. These patterns suggest that Gibrat's Law holds, as firm size does not diverge,- signifying that Despite being larger and more polluting, these firms did not grow faster economically. The regression results support that emissions reductions were due to mitigation at the firm level, rather than due to shifts in market share.

B.3 Entry dynamics

The local projections already showed that the decline in emissions among incumbent firms came from large, highly polluting firms. However, this analysis was only conducted from the perspective of incumbents relative to other incumbents at the start of the policy. A more extreme form of the relocation of market share could be the evolution of entrants. A substantial body of literature on creative destruction emphasizes the role of new entrants in driving the development of novel innovative products Aghion and Howitt (1992); Klette and Kortum (2004); Acemoglu et al. (2018), which is often accompanied by a subsequent expansion in their market share Haltiwanger et al. (2013, 2016). To verify that the observed trends are not the result of shifts in market share among entrants affecting both all emission-intense firms under the ETS equally, I compare the trajectories of entrants relative to those of incumbent firms.

I only classify firms newly appearing in the financial data as real entrants. Cases where existing firms enter the ETS through the acquisition of a polluting facility or through expanded regulatory coverage due to increased stringency of the policy are excluded from the analysis.[25] The rate of entry is 1.03% - extremely small relative to the common entry rate in the economy (Haltiwanger et al., 2013), suggesting that entry in itself plays a marginal role - however, this number does not specify the quality (in terms of clean technology) of these respective entrants.

Cross-sectional regressions demonstrate that entrants oftentimes start out smaller and more polluting than their incumbent counterparts. However, entrants are a special case of Gibrat's law (Haltiwanger et al., 2013) in terms of growth in size - thus potentially in terms of clean technology as well. To see whether they have a lower level of emission intensity relative to the incumbents and thus contribute positively to the aggregate reduction in emission intensity, I track them over time. More specifically, I track them for 7 years post entry. I do allow for attrition both at the level of incumbents and entrants. The regression is specified the following way:

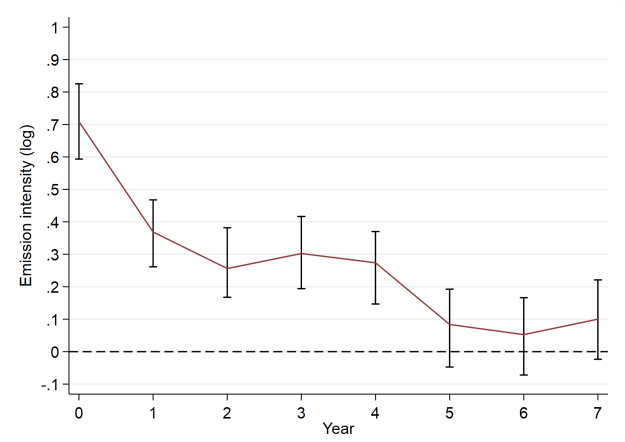
$$ln(ei_{i,t}) = \sum_{\tau=0}^{T=7} \alpha_\tau G_\tau D_i + \iota_{t,c} + \delta_{t,s} + \eta_{c,s} + \varepsilon_{i,t} \tag{B.2}$$

where $\ln(ei_{i,t})$ denotes the logarithm of emission intensity, G_τ measures years from the treatment event, and $\iota_{t,c}$ $\delta_{t,s}$, $\eta_{c,s}$ capture time-country, time-sector, and country-sector fixed effects respectively.

To generate a counterfactual group to which the evolution of entrants can be compared, I generate a placebo shock for a random group of 5% of incumbents in random years. Each incumbent can only receive 1 shock. This shock I consider a placebo entry decision and the evolution of the firms post shocks are then compared with the evolution of the entrants setting both decisions at $\tau = 0$. The variable D_i indicates whether the firm belongs to the actual entrants or the group of placebo entrants. This randomized placebo-

[25] The latter is captured in the following subsection - a direct comparison between the ETS-firms and the rest of the industry.

Figure B.3: Emission intensity of entrants



Note: Emission intensity of entrants compared to incumbents. A sample of placebo entrants is drawn from the pool of incumbents and tracked over the same seven-year period. This procedure is bootstrapped 1000 times (Young, 2019). The vertical intervals represent the 5th and 95th percentiles of the bootstrapped coefficients.

creation follows Young (2019) and is repeated 1000 times.

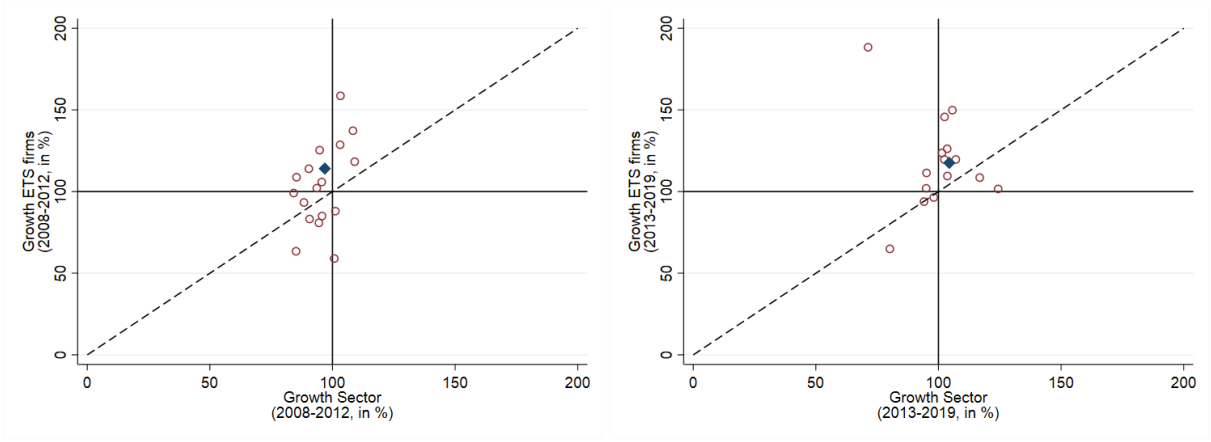
The results in Figure B.3 show that entrants initially exhibit significantly higher emission intensity than the incumbent comparison group. It is only around year 5 post-entry that their emission intensity converges, eliminating any structural difference between the two groups. This suggests that entrants do not actively contribute to the decline in aggregate emission intensity; at best, they pollute at similar levels as incumbents.

B.4 Relocation towards non-ETS

Finally, I investigate potential patterns of economic activity relocation between regulated (ETS-covered) and non-regulated firms within each nace 2-digit sector, I compare the growth rate of value added for the two groups. Potential divergent growth rates may indicate shifting activity away from regulated sectors. I distinguish across two Phase II and Phase III given their differing regulatory stringency and macroeconomic contexts.

Figure B.4 provides a comparative analysis of value-added growth (deflated) for ETS-covered firms relative to their respective sectors during Phase II (2008 - 2012) and Phase III (2013 - 2019) of the EU Emissions Trading System (ETS). The 45-degree line denotes parity in performance: points above this line indicate that ETS firms have outperformed their broader industry peers. In both subfigures, a majority of the observations lie above this parity line, suggesting that ETS firms consistently achieved higher growth in value added compared to their sectors. As a final check, the aggregate indicator for the entire manufacturing economy, represented by the blue diamond, also lies above the parity line in both phases. Overall, the figure suggests that participation in the ETS does not appear to have hindered firm-level growth; thus, relocation towards the non-ETS sector has been minimal.

Figure B.4: Growth value added relative to the sector



(a) Growth Phase II

(b) Growth Phase III

Note: Comparative analyses of growth in value added (deflated) for the firms active in the ETS and the whole sector. the x-axes denotes the growth of the sector in percentages relative to their economic situation at the beginning of the phase - while the y-axis shows this for the aggregate of ETS-firms in the sector. The 45-degree line indicates an equal performance of the ETS-firms and the industry. If the observation is above this line, the ETS-firms have performed better than their industry. The blue dot indicates the aggregate manufacturing economy.

B.5 Decomposition of the extensive margin

As a final check of the minor importance of the contributions of firm entry and exit to aggregate productivity dynamics, I apply the decomposition by Melitz and Polanec (2015) on the aggregate level of clean technology of firms operating in the European Union. As documented later, I consider clean technology $a_i = q_i/e_i$. This approach enables me to disentangle the sources of clean technology growth into contributions from incumbent firms, resource reallocation, and the net impact of entrants and exiters.

Formally, denote the change in aggregate clean technology $\Delta A_t = A_t - A_{t-1}$ where $A_t = \sum_i \omega_{i,t} a_{i,t}$ and $\omega_{i,t}$ denotes firm i's market share in value added at time t, with $\sum_i \omega_{i,t} = 1$. The decomposition can be stated as:

$$\Delta A_t = \overline{\Delta A_{S,t}} + \Delta \text{Cov}_S(\omega_{i,t}, a_{i,t}) + \text{Entry}_t + \text{Exit}_t, \tag{B.3}$$

Here, $\overline{\Delta A_{S,t}} = \sum_{i \in S} \omega_{i,t-1}(a_{i,t} - a_{i,t-1})$ captures the unweighted average productivity growth of surviving firms S weighted by their initial market shares (which is oftentimes termed the "within" effect). The second term $\Delta \text{Cov}_S(\omega_{i,t}, a_{i,t})$ captures the reallocation of market shares towards other firms in the industry (which one would consider the "between" effect). Finally $\text{Entry}_t = \sum_{i \in E} \omega_{i,t}(a_{i,t} - \overline{A}_{S,t})$, where E denotes the number of entrants, measures the contribution of new entrants relative to survivors, and $\text{Exit}_t = \sum_{i \in X} \omega_{i,t-1}(\overline{A}_{S,t-1} - a_{i,t-1})$, where X denotes exiters, accounts for the gap in clean technology relative to survivors in the prior period. I consider the EU ETS as one industry and omit the firms as discussed in Appendix A.

Table B.6 reports the decomposition of changes in aggregate clean technology (ΔA_t) for firms in the

Table B.6: Decomposition aggregate Clean technology growth

year	ΔA_t	$\Delta A_{S,t}$	$\Delta \text{Cov}_S(\omega_{i,t}, a_{i,t})$	Entry_t	Exit_t
2009	.0079	.0064	.0006	.0007	.0002
2010	.0057	.0119	-.0058	.0001	-.0002
2011	.0116	.0024	.0094	-.0004	.0002
2012	.0131	.0032	.0105	-.0009	.0001
2013	-.0176	-.0042	-.0128	-.0002	.0002
2014	.0088	.0071	.0018	.0003	.0001
2015	.0102	.0009	.01	-.0005	0
2016	.0025	.0047	-.002	-.0005	-.0008
2017	-.009	.0081	-.016	-.0004	0
2018	.0145	-.0003	.0129	.0008	-.0002
2019	-.0031	.0035	-.0053	-.0002	.0001
Mean	.0043	.0040	.0004	-.0001	-.0000

Note: Decomposition aggregate Clean technology growth following a dynamic Olley-Pakes decomposition by Melitz and Polanec (2015) for the years 2009 to 2019. The columns are ordered based on Equation B.3.

EU ETS over the period 2009 to 2019. The results reveal substantial inter-annual heterogeneity in both magnitude and sign across components, making it difficult to interpret and pinpoint the exact driving force behind the . Years such as 2012 are characterized by pronounced positive growth, driven jointly by strong within-firm improvements and favorable reallocation dynamics, whereas years like 2013 and 2019 exhibit sharp declines, largely attributable to negative reallocation effects.

On average, within-firm productivity growth is the main driver of aggregate clean technology change, while entry and exit effects contribute only modestly. The prominence of the within-firm reallocation term indicates that firm-specific, idiosyncratic decisions play a key role. The mean reallocation term is slightly positive, reinforcing the idea that these individual decisions are more influential than broader sector-level reallocation. This modest positive effect of reallocation is consistent with the direct impact of taxation on the distribution of firms. This tends to have regressive consequences on the clean technology margin as more polluting firms are more hit.

This decomposition differs slightly from the approaches in the literature that decompose pollutant emissions (Barrows and Ollivier, 2018; Rottner and von Graevenitz, 2024), which use an adaptation of the method proposed by Foster et al. (2008). In contrast, I follow the decomposition in Melitz and Polanec (2015), as it aligns more directly with the structure of the structural model and has an easier interpretation. Nonetheless, the results are consistent with the common finding that most aggregate technical change occurs within firms (Najjar and Cherniwchan, 2021; Shapiro and Walker, 2018; Barrows and Ollivier, 2018; Rottner and von Graevenitz, 2024). Since the decomposition is carried out in logs, it is equivalent to expressing the underlying variable either as clean technology adoption or, inversely, as emission intensity.

C. The Structural Model

C.1 A stylized model

This section develops a stylized model to provide theoretical structure for the intuition that the domestic government may have an incentive to transition gradually toward a fully decarbonized state.

Their incentive to choose a path depends on 3 key moments: the distribution of firms, the differential in taxation, and finally the disutility in the welfare function of current emissions carbon rather than the flow.

C.1.1 Setup

There are two time periods, 1 and 2, the 2nd period is discounted with a factor β. Consider three types of firms indexed by $i \in \{L, M, H\}$, each endowed with a level of clean technology denoted $a \in \{1, 10, \infty\}$, respectively. Firms are subject to a carbon tax ζ on their emissions of CO2. Profits for firm i are given by:

$$\pi_i = p_i q_i - \zeta e_i \tag{C.4}$$

where q denotes output, p is the product price, and e_i represents emissions. Production has a baseline unit cost of $c > 0$ equal to all firms. Emissions are modeled as a function of output and technology: $e_i = q_i / a_i$, with the cleanest firms ($a = \infty$) generating zero emissions. A carbon tax ζ applies to emissions. Under perfect competition, the price of variety i in period t is:

$$p_i = c + \frac{\zeta}{a_i} \tag{C.5}$$

where $\zeta / a_H = 0$.

The total number of firms is normalized to one: $N = \sum_i \mu_i = 1$ where μ_i defines the distribution of firms (or market share equivalently). Aggregate profits, output, and emissions are defined as $\Pi = \sum_i \mu_i \pi_i$, $Q = \sum_i \mu_i q_i$, and $E = \sum_i \mu_i e_i$, respectively. Consumers have Cobb-Douglas preferences, normalized to 1, with the distribution of firms μ_i. Income is normalized to 1, so $q_i = \mu_i / p_i$ and $\sum_i p_i q_i = 1$. Firms may upgrade their technology by incurring a fixed cost k, which enables them to move one step upward along a clean technology ladder (e.g., from L to M, or from M to H). To facilitate notation, this cost is assumed constant across technology levels.

Consumers allocate a fixed budget share to each good (Cobb-Douglas). Demand for each good can be written as:

$$q_i = \frac{1}{p_i} \tag{C.6}$$

Each period the government has access to two policy instruments: a low carbon tax, ζ^L, and a high carbon tax, ζ^H with $\zeta^H > \zeta^L$. A low tax induces only the most polluting firms (L-types) to upgrade, while a high tax induces both L- and M-type firms to adopt cleaner technology, but at a cost of reduced profits.

C.1.2 timing

Suppose the government aims to achieve a fully decarbonized (net-zero) economy within two periods. While doing so, it aims to balance the environmental objective with the preservation of domestic output. They can select two alternative policy pathways to get to net-zero $\theta \in \{A, B\}$: Path A is the low-high type, where the government gives time for the lower firms to adopt the clean technology. The second, path B, is the more urgent approach and reduces emissions the fastest at the cost of more output losses for the lowest type firms. Formally: $A = \{\zeta_1^D, \zeta_2^U\}$ and $B = \{\zeta_1^U, \zeta_2^U\}$

The government's objective function is:

$$\max_{\theta \in \{A,B\}} \sum_{t}^{2} \beta^{t-1} \left[Q_t(\theta) - \lambda E_t(\theta) \right] \quad \text{s.t. } E_2 = 0, \tag{C.7}$$

where $\theta \in \{A, B\}$ denotes the 2 potential policy paths, $\beta \in (0, 1]$ is the intertemporal discount factor, and $\lambda > 0$ reflects the government's relative weight on emissions reduction. Both aggregate emissions and aggregate demand depend on the chosen policy path θ.

path A: Low - High Because of the low tax in period 1, only the lowest number of firms adopt, while the rest do nothing. At period 1, Prices (as functions of ζ_A^1) are

$$p_{M,1}(\zeta_{A,1}) = c + \frac{\zeta_{A,1}}{10}, \qquad p_{H,1}(\zeta_{A,1}) = c.$$

Aggregate output and emissions are

$$\begin{aligned} Q_{A,1} &= [(\mu_L + \mu_M) q_M + \mu_H q_H] \\ &= \frac{\mu_L + \mu_M}{c + \zeta_{A,1}/10} + \frac{\mu_H}{c} \end{aligned} \tag{C.8}$$

$$\begin{aligned} E_{A,1} &= [(\mu_L + \mu_M) e_M + \mu_H e_H] \\ &= \frac{\mu_L + \mu_M}{10c + \zeta_{A,1}} \end{aligned} \tag{C.9}$$

In period 2, all remaining firms will obtain the high type: and total emissions $E_{2,A} = 0$ and aggregate output equals $Q_{A,2} = \frac{\mu_L + \mu_M + \mu_H}{c} = \frac{1}{c}$

Path B: High - High Under Path B, a high tax $\zeta_{B,1}$ in period 1 induces both L- and M-type firms to upgrade. Prices (as functions of $\zeta_{B,1}$):

$$p_{M,1}(\zeta_{B,1}) = c + \frac{\zeta_{B,1}}{10}, \qquad p_{H,1}(\zeta_{B,1}) = c.$$

Aggregate output and emissions are

$$\begin{aligned} Q_{B,1} &= [\mu_L q_M + (\mu_M + \mu_H) q_H] \\ &= \frac{\mu_L}{c + \zeta_{B,1}/10} + \frac{\mu_M + \mu_H}{c} \end{aligned} \tag{C.10}$$

$$\begin{aligned} E_{B,1} &= [(\mu_L)e_M + (\mu_M + \mu_H)e_H] \\ &= \frac{\mu_L}{10c + \zeta_{B,1}} \end{aligned} \tag{C.11}$$

In period 2, all firms have upgraded, yielding $E_{2,B} = 0$ and aggregate output equals $Q_{A,2} = Q_{B,2} = \frac{\mu_L + \mu_M + \mu_H}{c} = \frac{1}{c}$

C.1.3 Comparing paths

The selection for either path depends on the amount of emitted emissions and the produced output. The difference in emissions between the two paths is:

Differences (B minus A): For emissions

$$\Delta E_1 = E_{B,1} - E_{A,1} = \frac{\mu_L}{10c + \zeta_{B,1}} - \frac{\mu_L + \mu_M}{10c + \zeta_{A,1}}$$

Since $\zeta_{B,1} > \zeta_{A,1}$, Path B strictly reduces period-1 emissions more than Path A.

For quantities, the sign is generally ambiguous when $\zeta_{B,1} \neq \zeta_{A,1}$:

$$\begin{aligned} \Delta Q_1 &= Q_{B,1} + Q_{A,1} \\ &= \mu_M \left(\tfrac{1}{c} - \tfrac{1}{c + \zeta_{A,1}/10} \right) + \mu_L \left(\tfrac{1}{c + \zeta_{B,1}/10} - \tfrac{1}{c + \zeta_{A,1}/10} \right) \end{aligned} \tag{C.12}$$

The sign of ΔQ_1 - whether path B is preferred over path A - depends on the mass of firms with μ_M pushing the sign upwards and μ_l pushing it downwards. if the differences between prices are large $\zeta_B - \zeta_A$, then path A is preferred.

The welfare equation for time 1 in both paths can be written as:

$$\begin{aligned} W_{A,1} &= \frac{\mu_L + \mu_M}{c + \zeta_{A,1}/10} + \frac{\mu_H}{c} - \lambda \frac{\mu_L + \mu_M}{10c + \zeta_{A,1}}. \\ W_{B,1} &= \frac{\mu_L}{c + \zeta_{B,1}/10} + \frac{\mu_M + \mu_H}{c} - \lambda \frac{\mu_L}{10c + \zeta_{B,1}}. \end{aligned} \tag{C.13}$$

Thus their difference

$$\begin{aligned} \Delta W &= W_B - W_A \\ &= \left[\frac{\mu_L}{c + \zeta_{B,1}/10} + \frac{\mu_M}{c} - \frac{\mu_L + \mu_M}{c + \zeta_{A,1}/10} \right] - \lambda \left[\frac{\mu_L}{10c + \zeta_{B,1}} - \frac{\mu_L + \mu_M}{10c + \zeta_{A,1}} \right]. \end{aligned} \tag{C.14}$$

Intuitively, path B is favored by (i) the number of polluting firms is extremely small (μ_L), (ii) or when the government has a high level of emissions preference (λ is large). In that case where emissions outweigh the positive effect of quantities entirely, path B dominates. In contrast, the government prefers path A over path B when (i) the price to convince the number of middle firms to adopt novel technologies is very high (ζ_B); (ii) there are many low-type firms (μ_L large) or (iii) emissions concerns are considered less important.

Formally for a λ sufficiently small (< 10):

- ΔW rises with μ_L when $\lambda < 10$, and falls when $\lambda > 10$.

$$\frac{\partial \Delta W}{\partial \mu_L} = (10 - \lambda)\left[\frac{1}{10c + \zeta_B} - \frac{1}{10c + \zeta_A}\right],$$

- ΔW rises with μ_M if $\zeta_A + \lambda c > 0$; and falls if $\zeta_A + \lambda c < 0$.

$$\frac{\partial \Delta W}{\partial \mu_M} = \frac{\zeta_A + \lambda c}{c(10c + \zeta_A)},$$

- ΔW falls with ζ_B if $\lambda < 10$, and rises if $\lambda > 10$.

$$\frac{\partial \Delta W}{\partial \zeta_B} = \mu_L \frac{\lambda - 10}{(10c + \zeta_B)^2},$$

- ΔW rises with ζ_A if $\lambda < 10$, and falls if $\lambda > 10$.

$$\frac{\partial \Delta W}{\partial \zeta_A} = (\mu_L + \mu_M) \frac{10 - \lambda}{(10c + \zeta_A)^2},$$

C.2 Notation

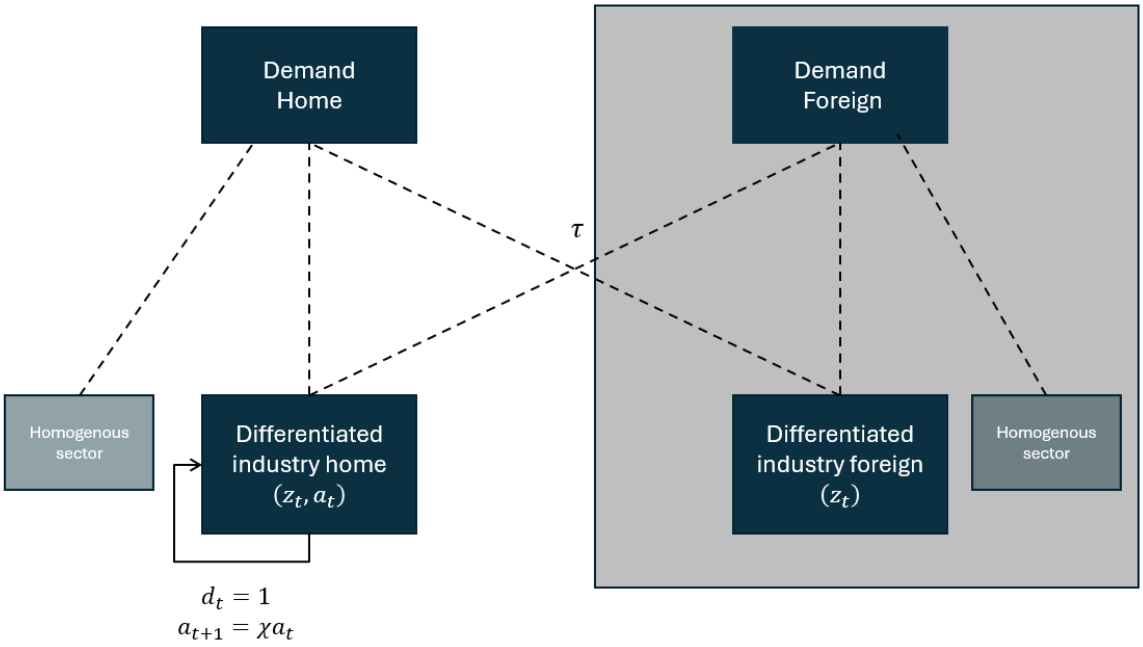
Table C.7: Notation of variables in the model

Parameter	Explanation
ζ	Price of a ton of carbon
κ	Allocation of free certificates
a	level of clean / abatement technology
z	level of productivity
z^T	transitory component productivity
ξ,σ_ξ	idiosyncratic productivity shock and stdv.
p	price of the differentiated good
q	quantity of the differentiated good
c	marginal cost $c \equiv (1 + z\zeta/a) \geq 1$
τ	trade cost
e	emissions
π	profits
β	Discount factor
η	Elasticity of substitution demand
d, k	Decision to invest and adopt novel technology ($d(\chi) = 1$) and its cost
χ	step-size increase technology adoption
k^F,k^C	Fixed and convex component adoption cost
γ	Evolution of adoption cost k, with $\gamma \in (0, 1)$
ω	Evolution of productivity
P	Aggregate price index differentiated good
Q	Aggregate Quantity differentiated good
I	Aggregate Income allocated to differentiated good
E	Aggregate Emissions emitted by production differentiated good
W	Aggregate wage
Υ, R	Composite of all aggregate forces
F	motion parameter of distribution of firms
μ	distribution of firms
β_G	Government Discount factor

Note: Explanation of the variables used in the structural model. Aggregate variables are denoted in capital letters. The time-dimension t is omitted. If the variable is indexed with d or x it denotes whether it is produced for domestic market or the export market. If it is indexed with h or f it states whether it is at home or at the foreign economy. Whenever the variables are indexed with an aggregate D or S they indicate whether they are used for demand or supply respectively. For instance Q_h^D denotes the aggregate demand of the differentiated good in the home economy.

Table C.7 provides a short summary of the variables used in the structural model in the text. Note that this table does not include any of the variables in the model variation used in the appendix, nor in the regressions and some of the letters used in the appendix have different meanings depending on the context.

Figure C.5: Directed Acyclic Graph (DAG)



Note: visual overview of the static building blocks of the structural model.

To give a quick visual overview of the structure of the structural model, figure Figure C.5 illustrates the dependencies between the different agents. The 'passive' foreign country is depicted in grey. A second homogeneous sector is included to fix prices, though it is normalized away.

Note that a large part of the model plays out on the dynamic front. The evolution of aggregate prices - both in the input market as in the output market - as well as the endogenous adoption decision.

C.3 Solving the model

The model is solved using a nested numerical methods approach, where I first solve for the aggregate price index given an exogenously determined carbon path and end of state value function and V_T. The model updates such that it solves for all criteria. After the model is solved, it is estimated using a MLE strategy, which aims to find the adoption cost and evolution to minimize the difference with the observed data.

More specifically; first, steady state 1 is solved for the aggregate price. This is a simple analytical procedure with the distributions of the 2 productivity state variables known. Since the carbon policies are set to zero, making firms unable to deviate from their current state. I set the value of the ex-post value function to be equal to the one of the ex-ante distribution - however $a = a^{max}$ for all firms.

Here onward, I start with an iterative procedure to minimize the maximum likelihood (Equation 20). The model consists of a nested iterative loop with 2 loops. I start from the inner loop and diverge.

1. Inner loop: A random path of carbon prices is drawn for the whole period. Given (i) this path of carbon prices, (ii) the exogenously set carbon cap and (iii) the ex-ante determined aggregate distribution from the first steady state, I search for the yearly aggregate price indices. in the domestic market P_h. As the distribution of firms is set and aggregate demand is stable, the market in the foreign market clears automatically given P_h - which in turn gives the foreign aggregate price index too.

2. Outer loop: given the aggregate prices and carbon price, the sum of carbon emissions is computed. This loop minimizes the differential between the demanded sum of carbon emissions and the exogenously determined cap. The price of carbon is updated.

3. After the model is solved, I estimate it using a simple maximum likelihood estimator with a Nelder-Mead simplex algorithm. The maximum likelihood procedure aims to find the values of $\{k^F, k^Q, \gamma\}$

The model consists of 2 state variables: There are 5 grid values for productivity and 161 values of clean technology. The criteria to solve the model for each iteration can be summarized as the maximum relative error between the new and old values over the time steps:

$$\text{crit} = \max_{1 \leq t \leq T} \left| \frac{x^{new}(t) - x^{old}(t)}{x^{old}(t)} \right| \tag{C.15}$$

where x equals the aggregate price index in the first iteration (P_h) and the aggregate domestic demand for emissions in the second iterative procedure (E_h^D). To update the carbon price - aggregate prices, I apply a simple binary search method per-time-step t.

C.4 Decompositions

C.4.1 Decomposing aggregate price index

To decompose the shifts of the aggregate price index in the home country ΔP_h, I start by defining 2 new indices: P_h^A and P_h^* denote the actual price index and the optimal price index where there are no carbon policies in place, respectively. For readability and space considerations, I deviate from the notation in the

text and write P^0, P^λ as the aggregate price indices given each decision (respectively none and for each adoption decision $\lambda \in \{1.05, 1.10, 1.20\}$).

$$\begin{aligned}\Delta P_h &= P^A - P^* \\ &= \underbrace{P_h^A - P_h^0}_{\text{Decisions}} + \underbrace{P_h^0 - P_h^*}_{\text{taxation}} \\ &= \underbrace{P_h^0 - P_h^*}_{\text{taxation}} + \sum_\lambda \underbrace{P_h^\lambda - P_h^0}_{\text{adoption } \lambda} \qquad \text{where } P_h^A - P_h^0 \approx \sum_\lambda P_h^\lambda - P_h^0\end{aligned} \tag{C.16}$$

This decomposition holds when each adoption decision is independent and additive or when the impact of either decision on its respective price index is sufficiently small (i.e. $P_h^\lambda \approx P_h^0$), then the decomposition is unbiased. Since the decisions are mutually exclusive by assumption, then the support of $\mu^\lambda - \mu^0$ is disjoint, thus $\mu_h^A = \mu_h^0 + \sum_\lambda (\mu_h^\lambda - \mu_h^0)$.

To see how this generates an unbiased result in the aggregate price index, I start by taking the logarithm of each of the aggregate price indices. Let

$$f(z,a) := p_{hd}^{1-\eta}, \qquad g(z,a) := p_{fx}^{1-\eta} = Df(z,a), \qquad D := \tau^{1-\eta},$$

and define $\widetilde{\mu} = \mu_h + D\mu_f$. Then for each scenario $x \in \{A, 0, \lambda, *\}$,

$$\begin{aligned}ln(P_{h,t}^X) &= \frac{1}{1-\eta} ln\left(\int_{z*a} f(z,a) d\mu_h^X + \int_{z*a} g(z,a) d\mu_f^X\right) \\ &= \frac{1}{1-\eta} ln\left(\int_{z*a} f(z,a) d\mu_h^X + D\int_{z*a} f(z,a) d\mu_f^X\right) \\ &= \frac{1}{1-\eta} ln\left(\int_{z*a} f(z,a) d[\mu_h^X + D\mu_f^X]\right) = \frac{1}{1-\eta} ln\left(\int_{z*a} f(z,a) d[\tilde{\mu}^X]\right)\end{aligned} \tag{C.17}$$

If each adoption effect is small (i.e. $|\delta_\lambda| \ll 1$), use $\ln(1+x) \approx x$. Then

$$\ln P_{h,t}^A - \ln P_{h,t}^0 \approx \frac{1}{1-\eta} \sum_\lambda \delta_\lambda \quad \text{and} \quad \ln P_{h,t}^\lambda - \ln P_{h,t}^0 \approx \frac{1}{1-\eta} \delta_\lambda,$$

so up to a first order approximation the equation can be stated as:

$$\ln P_{h,t}^A - \ln P_{h,t}^0 \;\approx\; \sum_\lambda \left(\ln P_{h,t}^\lambda - \ln P_{h,t}^0\right).$$

C.4.2 Decomposing aggregate emissions

Rewrite where E^A and E^* denote actual emissions and the business as usual scenario without carbon policies in place.Let E^0 denote emissions under taxation alone, and E^λ denote emissions under individual

adoption/relocation decisions λ. Aggregate can be decomposed as follows:

$$\begin{aligned}\Delta E_h &:= E^A - E^* \\ &= \underbrace{E^A - E^0}_{\text{decisions}} + \underbrace{E^0 - E^*}_{\text{taxation}} \\ &= \underbrace{E^0 - E^*}_{\text{taxation}} + \sum_\lambda \underbrace{E^\lambda - E^0}_{\text{adoption } \lambda} \qquad \text{(where: } E^A - E^0 \approx \sum_\lambda (E^\lambda - E^0)\text{).}\end{aligned} \tag{C.18}$$

Conditional on aggregate prices: This decomposition holds exactly when the impact of the distribution on domestic emissions is linear, thus the aggregate price index is fixed. More specifically, I define this counterfactual using the actual aggregate price P^A:

$$\Delta_h(P_h^A, P_f^A) = \underbrace{E^0(P_h^A, P_f^A) - E^*(P_h^A, P_f^A)}_{\text{taxation}} + \sum_\lambda \underbrace{E^\lambda(P_h^A, P_f^A) - E^0(P_h^A, P_f^A)}_{\text{adoption } \lambda} \tag{C.19}$$

$$E_h^A = C\int_{z,a} f(z,a,\zeta)\, d\mu_h^A(z,a), \quad E_h^0 = C\int_{z,a} f(z,a,\zeta)\, d\mu_h^0(z,a), \quad E_h^\lambda = C\int_{z,a} f(z,a,\zeta)\, d\mu_h^\lambda(z,a) \tag{C.20}$$

where $C = \left(\frac{\eta}{\eta-1}W\right)^{-\eta}(P_h^A)^{\eta-1}$. and $f(z,a,\zeta) = \left(\frac{1}{z} + \frac{\zeta}{a}\right)^{-\eta}\frac{1}{a}$

Since both decisions are mutually exclusive, the supports of $\mu^\lambda - \mu^0$ are disjoint, giving $\mu^A = \mu^0 + \sum_\lambda(\mu^\lambda - \mu^0)$. .Using the linearity of the integral operator, aggregate emissions decompose exactly as:

$$\begin{aligned}E_h^A - E_h^0 &= C\left[\int f\, d\mu_h^A - \int f\, d\mu_h^0\right] \\ &= C\int f\, d\left[\mu_h^A - \mu_h^0\right] \\ &= C\sum_\lambda \int f\, d\left[\mu_h^\lambda - \mu_h^0\right] \\ &= \sum_\lambda \left(E_h^\lambda - E_h^0\right)\end{aligned} \tag{C.21}$$

Fluctuating aggregate prices: When allowing the aggregate price index to fluctuate, the decomposition still holds approximately to first order, because emissions enter linearly in the integral. Denoting the counterfactual prices by P_h^λ, P_f^λ, we have:

$$E_h^A(P_h^A, P_f^A) - E_h^0(P_h^0, P_f^0) \approx \underbrace{E^0(P_h^0, P_f^0) - E^*(P_h^*, P_f^*)}_{\text{taxation}} + \sum_\lambda \underbrace{E^\lambda(P_h^\lambda, P_f^\lambda) - E^0(P_h^0, P_f^0)}_{\text{adoption } \lambda}. \tag{C.22}$$

C.4.3 Other aggregate variables

All other aggregate variables can be decomposed using a similar decomposition strategy.

C.5 Implications of a carbon policy

This section distinguishes between the direct and indirect effects of a unilateral carbon policy. It considers the effect through the domestic price index and the foreign price index separately. Note that the ter "indirect" in this setting applies to the general equilibrium effect. It still the intratemporal indirect effects through the aggregate price indices. I consider the distribution of firms to stay constant. How the aggregate results are shaped by shifts in the distribution is explained in subsection C.4.

C.5.1 Aggregate (domestic) price index

I start by rephrasing the aggregate price index in the domestic market.

$$P_h = \left[\int_{z*a} p_{hd}(z,a)^{1-\eta} d\mu_h(z,a) + \int_{z*a} p_{fx}(z)^{1-\eta} d\mu_f(z,a) \right]^{\frac{1}{1-\eta}} \tag{C.23}$$

Filling in the functional forms for the idiosyncratic domestic prices $p_{hd} = \eta/(\eta-1)W(1/z+\zeta/a)$ and foreign prices $p_{fx} = \eta/(\eta-1)\tau W(1/z)$. The prior have internalized the implications of a carbon policy on their clean technology margin, while the others have not. In contrast the foreign firms have to pay an additional iceberg tax τ.

$$\begin{aligned} P_h &= \left[\left(\frac{\eta}{\eta-1} W \right)^{1-\eta} \int_{z*a} \left(\frac{1}{z} + \frac{\zeta}{a} \right)^{1-\eta} d\mu_h(z,a) + \left(\frac{\eta}{\eta-1} W \right)^{1-\eta} \tau^{1-\eta} \int_{z*a} z^{\eta-1} d\mu_f(z) \right]^{\frac{1}{1-\eta}} \\ P_h &= \left(\frac{\eta}{\eta-1} W \right) \left[\int_{z*a} \left(\frac{1}{z} + \frac{\zeta}{a} \right)^{1-\eta} d\mu_h(z,a) + \tau^{1-\eta} \int_{z*a} z^{\eta-1} d\mu_f(z) \right]^{\frac{1}{1-\eta}} \end{aligned} \tag{C.24}$$

The derivative w.r.t. ζ equals:

$$\frac{\partial P_h}{\partial \zeta} = \left(\frac{\eta}{\eta-1} W \right)^{1-\eta} P_h^{\eta} \int_{z*a} \left(\frac{1}{z} + \frac{\zeta}{a} \right)^{-\eta} \frac{1}{a} d\mu_h(z,a) > 0 \tag{C.25}$$

Consequently, a carbon policy affects the real prices positively.

Similarly, the foreign price index is affected similarly with the derivative w.r.t ζ being:

$$\frac{\partial P_f}{\partial \zeta} = \left(\frac{\eta}{\eta-1} \tau W \right)^{1-\eta} P_f^{\eta} \int_{z*a} \left(\frac{1}{z} + \frac{\zeta}{a} \right)^{-\eta} \frac{1}{a} d\mu_h(z,a) > 0 \tag{C.26}$$

C.5.2 Aggregate demand

The CES aggregator is:

$$Q_h = \left[\int q_{hd}(z,a)^{\frac{\eta-1}{\eta}} d\mu_h + \int q_{fx}(z)^{\frac{\eta-1}{\eta}} d\mu_f \right]^{\frac{\eta}{\eta-1}} \tag{C.27}$$

Where the term $q_{hd}(z,a)$ denotes the domestic quantity for each variety (z,a), while $q_{fx}(z)$ represents the foreign quantity for variety z. The parameter $\eta > 1$ corresponds to the elasticity of substitution. Finally, $d\mu_h$ and $d\mu_f$ are the measures over domestic and foreign varieties, respectively.

Using the chain rule, the total derivative of Q_h w.r.t. ζ:

$$\frac{\partial Q_h}{\partial \zeta} = \underbrace{\int_{z,a} \frac{\partial Q_h}{\partial q_{hd}} \frac{\partial q_{hd}}{\partial \zeta} d\mu_h}_{\text{Domestic varieties, intensive margin (-)}} + \underbrace{\int_{z} \frac{\partial Q_h}{\partial q_{fx}} \frac{\partial q_{fx}}{\partial \zeta} d\mu_f}_{\text{Foreign varieties, trade effect (+)}} + \underbrace{\frac{\partial Q_h}{\partial P_h} \frac{\partial P_h}{\partial \zeta}}_{\text{Aggregate price effect } P_h \text{ (+)}} \tag{C.28}$$

The first term captures the direct effect on domestic quantities produced, while the second term reflects the effect on imports of foreign quantities. This is zero. . The third term represents the response of the CES aggregator to changes in the price index.

(a) Domestic varieties: intensive margin The first term can be further decomposed into

$$\int_{z,a} \frac{\partial Q_h}{\partial q_{hd}} \frac{\partial q_{hd}}{\partial \zeta} d\mu_h = \int_{z,a} \underbrace{Q_h^{1/\eta} q_{hd}^{-1/\eta}}_{\partial Q_h / \partial q_{hd}} \underbrace{\left[-\eta \frac{q_{hd}}{p_{hd}} \frac{\partial p_{hd}}{\partial \zeta} + (\eta - 1) \frac{q_{hd}}{P_h} \frac{\partial P_h}{\partial \zeta} \right] d\mu_h}_{\partial q_{hd} / \partial \zeta} \tag{C.29}$$

The sign of the term depends entirely by the bracketed term: The first part in this bracket can be read as the direct "intensive-margin" effect of ζ on domestic prices, which is positive. The second term yields the negative indirect effect via the aggregate price index P_h. Consequently, it depends on the relative strength of either. Intuitively, on average all prices rise as ζ increases, yet firms with a high level of clean technology are less affected by this increase - to the point where some firms even face positive effects as domestic demand substitutes towards their cleaner products.

Filling in the functional forms for the price and the derivative for the aggregate price index. The equation becomes

$$\int_{z,a} \frac{\partial Q_h}{\partial q_{hd}} \frac{\partial q_{hd}}{\partial \zeta} d\mu_h = -\eta Q_h^{1/\eta} \int_{z*a} q_{hd}^{\frac{\eta-1}{\eta}} \frac{1}{p_{hd}} \frac{\eta}{\eta - 1} \frac{W}{a} d\mu_h + (\eta - 1) \frac{Q_h^{-1/\eta}}{P_h} \frac{\partial P_h}{\partial \zeta} \int_{z*a} q_{hd}^{\frac{\eta-1}{\eta}} d\mu_h(z,a) \tag{C.30}$$

(b) Foreign varieties: trade effect The 2nd component of the total derivative of Q_h w.r.t. ζ, the foreign varieties effect, yields the partial derivative of Q_h w.r.t q_{fx} as well as the partial derivative of q_{fx} w.r.t. ζ. The derivatives for either can be written as:

$$\frac{\partial Q_h}{\partial q_{fx}} = Q_h^{1/\eta} q_{fx}^{-1/\eta}, \quad \frac{\partial q_{fx}}{\partial \zeta} = (\eta - 1) \frac{q_{fx}}{P_h} \frac{\partial P_h}{\partial \zeta} \tag{C.31}$$

Since the foreign varieties are only affected indirectly through the aggregate price channel, the effect ζ on individual demand simplifies a whole lot and can only be read as positive. Intuitively, as domestic goods are getting more expensive, domestic demand is going to shift towards less expensive alternatives, ergo foreign varieties.

Consequently, as a whole, the trade effect can be reformulated as:

$$\int_{z} \frac{\partial Q_h}{\partial q_{hd}} \frac{\partial q_{fx}}{\partial \zeta} d\mu_f = (\eta - 1) \frac{Q_h^{1/\eta}}{P_h} \frac{\partial P_h}{\partial \zeta} \int_{z} q_{fx}^{(\eta-1)/\eta} d\mu_f > 0 \tag{C.32}$$

(c) direct index effect The final term reads the response of the CES aggregator the direct changes in the price index.

$$\frac{\partial Q_h}{\partial P_h}\frac{\partial P_h}{\partial \zeta} = (\eta - 1)\frac{Q_h}{P_h}\frac{\partial P_h}{\partial \zeta} > 0 \tag{C.33}$$

Under the usual economic assumptions, this is also strictly positive.

(d) Compact form Because of the budget clearing assumption and the CES structure of demand, $Q_h = P_h^{-1}$. The total effect of ζ can be rewritten as:

$$\frac{\partial Q_h}{\partial \zeta} = -P_h^{-2}\frac{\partial P_h}{\partial_\zeta} < 0 \tag{C.34}$$

And is strictly negative. It is therefore needed that the first component - the domestic varieties component which is analytically ambiguous in its sign - outweighs the other two components to achieve a total negative effect on aggregate demand.

C.5.3 Domestic Emissions

All derivatives here are conducted through the assumptions that firms are price takers in the carbon market and lowering the carbon cap yields a rise in the price of the carbon cap. Here I demonstrate that given these assumptions, the relationship is monotone. I start by defining an implicit function of the arg max equation in the text:

$$F(\zeta, E_h^S) = E_h(\zeta) - E_h^S = 0 \tag{C.35}$$

Consequently, the derivative of ζ w.r.t. E_h^S equals

$$\frac{\partial \zeta}{\partial E_h^S} = -\frac{\partial F/\partial E_h^S}{\partial F/\partial \zeta} = \frac{1}{\partial E_h^D/\partial \zeta} < 0 \tag{C.36}$$

The fact that the sign of this equation is positive can be demonstrated through the direct effect of ζ on domestic emissions:

$$\frac{\partial E_h}{\partial \zeta} = \int_{z*a}\frac{1}{a}\frac{\partial q_{hd}}{\partial \zeta}d\mu(z,a) + \int_{z*a}\frac{1}{a}\frac{\partial q_{hx}}{\partial \zeta}d\mu(z,a) < 0 \tag{C.37}$$

Where the first derivative constitutes the output for domestic usage and the second the output used for exports. As discussed in the partial derivative of Q_h w.r.t. ζ there is no direct solution for $\partial q_{hd}/\partial \zeta$ yet the derivative has to be negative to sum up all other positive effects and obtain an aggregate negative effect. This applies too for $\partial q_{hx}/\partial \zeta$.

C.6 Evolution of the carbon cap

C.6.1 Evolution from to 2050

I constrain the carbon emissions cap to follow a quadratic function of the form $y = b_1x^2 + b_2x + b_3$, ensuring it passes through two specific points on the timeline: $(0, 100)$, representing the baseline year when emissions

are not yet targeted (year 0), and $(42,10)$, corresponding to the net zero target for 2050 (year 42). The first point $(0,100)$ determines $b_3 = 100$. The second point determines $b_2(b_1)$.

Solving for b at $x = 42$, gives:

$$\begin{aligned} 10 &= a(42)^2 + b(42) + 100 \\ -90 &= 1764a + 42b \\ b &= -\frac{90}{42} - 42a \end{aligned} \tag{C.38}$$

Thus the equation becomes a function of the first parameter: $y = b_1x^2 + x(-90/42 - 42b_1) + 100$.

Finally, I constrain the limits of the function such that:

$$y = \begin{cases} 10 & \text{if } y < 10 \\ y & \text{if } 10 \leq y \leq 100 \\ 100 & \text{if } y > 100 \end{cases}$$

C.6.2 Carbon forward guidance

Under carbon forward guidance where the government aims to maximize output rather than an emission-weighted welfare function (i.e. $\lambda = 0$), the cap is non-binding until the very last observation at time T. This generates one additional degree of freedom. There is an inequality constraint on the second to last observation of the cap $E^S_{h,T-1} > E^S_{h,0}$. Additionally $E^S_{h,t} = E_{h,0}$ for all $t < T$. The additional degree of freedom makes it possible to find the maximum level of b_1 for which this equation holds. Filling in for $(100,41)$ yields the following:

$$100 = b_1x^2 + x(-90/42 - 42b_1) + 100 \tag{C.39}$$

This makes $b_1 = -615/287 \approx -2.14$. Consequently, the cap holds for $b_1 < -2.14$.

For completion, the opposite can also be done with a government that focuses solely on the negative externalities of CO2 emissions. In that scenario, the government achieves net zero in the following period $(10,1)$. $b_1 > 615/287 \approx 2.14$. I consider this scenario as implausible. It is just added to provide more intuition on the bounds of b_1.

D. Structural Estimation

This section provides additional information on the construction of the parameters used to estimate the model: clean technology $a_{i,t}$, productivity $z_{i,t}$, which can be further decomposed in a persistent z_i^F and transitory $z_{i,t}^T$ component, and the decision variable $d_{i,t}$ on whether to adopt novel clean technologies. I then show the non-parametric distribution. Finally, I back out the trade cost τ based on the methods proposed by Head and Ries (2001).

D.1 construction of indicators

The estimation strategy of the model requires the creation of various variables outside the scope of the model. This section provides a detailed account of how these variables are generated and presents a series of robustness checks to validate the reliability of the constructed indicator.

I begin by constructing a value-added production function following the approach of Asker et al. (2014) to derive a firm-level measure of productivity. This productivity measure is subsequently decomposed into its persistent and transitory components. Based on the production function, I then infer a measure of clean technology, defined as the ratio of value added to emissions â effectively the inverse of emission intensity.

Next, I develop an indicator designed to capture structural improvements in clean technology. To ensure that this indicator truly reflects technological upgrading rather than mere output retraction, I examine whether (i) it is associated with increases in capital investment rather than only with higher value added. I further regress this indicator on the carbon price to assess whether the EU ETS has influenced the adoption decision as suggestive evidence that decisions are based on changes in the carbon policy. Finally, to evaluate whether such adoptions are of first-order economic significance, I compare the ex-ante characteristics of firms undertaking these investments with those of firms that exhibit the largest emission reductions, following the local projection framework introduced in the descriptive analysis. They are the same type of firms. Thus, those that adopt are also those that reduce their emissions mostly.

D.1.1 Firm productivity

To obtain the idiosyncratic productivity components, I follow Asker et al. (2014). The production for a profit-maximizing firm i at time t is restated and augmented. Where for the model I am not specific on the input mix used. Here I specifically state the usage of 2 inputs, labor l and capital k. The production function can be written as:

$$q_{it} = \min\{e^{z_{it}} f(k_{it}, l_{it}), e_{it} a_{it}\} \tag{D.40}$$

where he parameter of productivity z_{it} is measured as a residual of a production function $f(.)$ with value added q as an output and two inputs, capital k an labor l for each firm i in time t. Relative to the baseline model, in here, I am explicit in the input factors to obtain z. Clean technology a can still be seen as a convex combination of both inputs.

I obtain revenue productivity, as argued by Foster et al. (2008), since, as common in most of the literature, I do not observe prices and quantities separately. The definition of productivity in this specification

also includes differences in markups, adjustment costs and other firm-specific factors (Asker et al., 2014). The chosen functional form of the production function is Cobb-Douglas in its inputs $f(.)$, and Leontief in emissions e interacted with the level of clean technology a. I assume that the production function exhibits constant returns to scale (i.e. $\alpha_l + \alpha_k = 1$).

The firm faces a downward-sloping demand curve with constant elasticity:

$$q_{it} = B_{it} p_{i,t}^{-\eta}, \tag{D.41}$$

where $p_{i,t}$ are the firm-specific prices and η denotes the elasticity of substitution. The parameter $B_{i,t}$ captures additional general equilibrium effects, potential spillovers, and any idiosyncratic shocks affecting the firm's demand.

Combining the demand function with output yields the following value-added production function (Ackerberg et al., 2015):

$$va_{it} = \omega_{i,t} l_{i,t}^{\beta_l} k_{i,t}^{\beta_k}, \tag{D.42}$$

where $\omega_{i,t} = (e^{z_{it}})^{(}1 - 1/\eta)B_{i,t}^{\eta}$. The elasticity parameters are distortions from the actual parameters with $\beta_l = \alpha_l(1 - 1/\eta)$ for labor.

To recover the correct measure of productivity ($z_{i,t}$), I compute the expenditure share of labor. Under the assumption that there are no additional adjustment costs, the parameter β_l and β_k equal their respective expenditure share. Following Asker et al. (2014), I compute β_l as the median labor expenditure share across all firms in the sample of ETS firms:

$$\beta_l = \text{median}\left(\frac{Wl_{i,t}}{va_{i,t}}\right) \tag{D.43}$$

where $Wl_{i,t}$ denotes the cost of employees in the ORBIS database. For capital, I follow Bartelsman et al. (2013); Van Beveren (2012) among others and select tangible fixed capital.

Given the constant returns to scale assumption, the coefficient for capital β_k can be obtained: $\beta_k = (\eta - 1)/\eta - \beta_l$, where $\eta = 6$, consistent with the model. Finally, given these estimates, productivity can be backed out as:

$$ln(\omega_{it}) = ln(va_{i,t}) - \beta_l ln(l_{it}) - \beta_k ln(k_{it}) \tag{D.44}$$

The estimation results of the two output elasticities for the sample of ETS-firms are given in Table D.8.

D.1.2 Productivity decomposition

Baseline The only restriction set on the observed productivity z is the condition of it being a Markov process: $z_{it} = g(z_{it-1}) + z_{it}^T$. It can be further decomposed into a permanent and a transitory component, where all time variation is captured by the transitory component. I restrict the process for the transitory component further to be white noise process without any persistence.[26]

$$z_{i,t} = z_i^F + z_{i,t}^T, \qquad \text{with} \quad z_{i,t}^T \sim N(0, \sigma_z^2) \tag{D.45}$$

[26]This restrictions will be relaxed in stages.

where innovation to the transitory component is i.i.d..

Table D.8: Productivity estimation strategy

Productivity	β_k	β_l	Observations	No. of firms
Values	0.552	0.281	38.073	3493

Note: the table shows output elasticities of the factor share production function estimation based on Asker et al. (2014). The lower part demonstrates the outcome variable of ρ in Equation F.55. The standard deviation of the transitory component equals $\sigma(z_{it}^T) = 0.410$.

The regression in Equation F.55 makes no assumption on the parametric distribution of z_i^F. It does so for $z_{i,t}^T$, which is drawn from a normal distribution and i.i.d.

D.1.3 Clean technology margin

Clean technology Under the assumption that emissions are Leontief in the value-added production function, the abatement technology can be written as a ratio of the emissions and value added. Formally:

$$a_{it} = \frac{q_{it}}{e_{it}} \approx \frac{va_{it}}{e_{it}} = \frac{a_{it} p_{it} q_{it}}{q_{it}}, \tag{D.46}$$

where p_{it} are deflated at the country-sector level using Eurostat deflator data. Note that - akin to the production estimation due to the absence of firm-specific prices - this is a measure of revenue clean technology rather than clean technology Foster et al. (2008). The inverse of this measure is oftentimes called emission intensity - the degree of emissions used to produce one unit of output. I do not use this for the reason that it is numerically more stable in the empirical regressions and is more intuitive for the quality ladder design of clean technology used in the model.

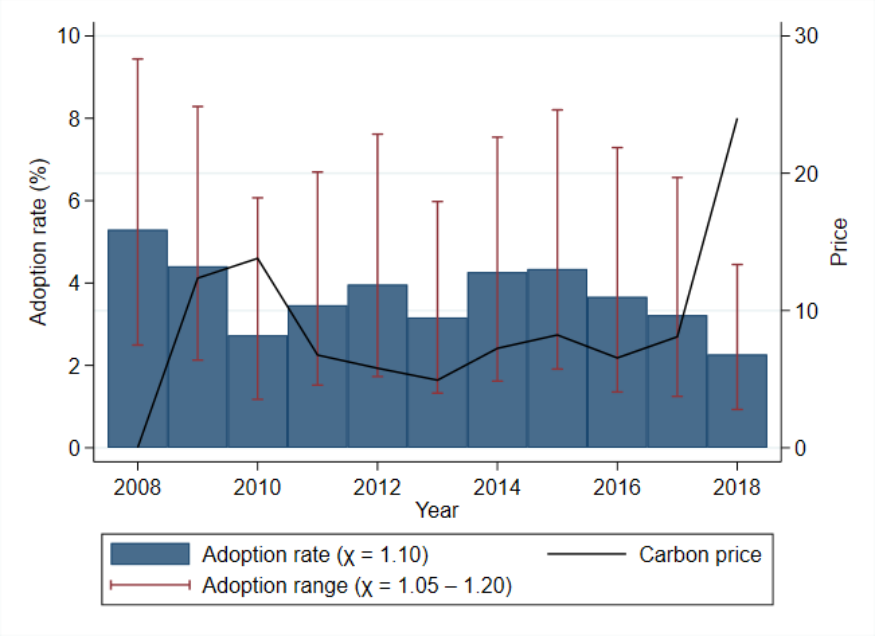
Clean technology adoption The decision to undertake emission abatement is both costly and infrequent, often characterized by lumpiness and persistence. While direct costs of abatement are unobservable, changes in the level of clean technology $a_{i,t}$ provide a measurable outcome through which successful abatement can be inferred. To capture these changes, I construct a novel binary indicator $d_t(\chi) \in \{0,1\}$ inspired by Bessen et al. (2019), which signals whether firm i engaged in abatement activity, with at least a magnitude χ, in year t. The indicator can be formalized as:

$$d_t(\chi) = \mathbb{1} \underbrace{\{\chi \tilde{a}_{i,t+1} \geq \tilde{a}_{i,t}\}}_{\text{Transitory state}} \underbrace{\left\{ \chi \frac{1}{B} \sum_{b=1}^{B=2} \tilde{a}_{i,t+b} \geq \frac{1}{W+1} \sum_{w=0}^{W=2} \tilde{a}_{i,t-w} \right\}}_{\text{structural state}} \tag{D.47}$$

where the indicator $d_t(\chi)$ equals 1 only if two conditions are met. First, there is a *transitory state* which captures a shock to clean technology at time t, exceeding a threshold $\chi \in \{1.05, 1.10, 1.20\}$ reflecting the

lumpiness of such increase.[27] Second, the *structural state* ensures the absorption of the firm into a new state that is persistently higher than the initial state, with average clean technology over the next two years at least $(\chi - 1)\%$ higher than the previous three-year average.

Figure D.6: Yearly adoption rate



Note: Yearly adoption of higher levels of clean technology. The blue bars denote the occurrences when Equation D.47 exceeds the threshold of $\chi > 1.10$. The red ranges show the result for 1.05 and 1.20 (the latter being the lower part of the range). The figure is constructed in such a way that they are mutually exclusive: If the adoption is more than 20% it is excluded from the earlier threshold values. In black, the carbon price is depicted in euros/tonne.

While this indicators highlights structural shocks in clean technology at the firm level, it does not define the driving nature behind this structural change. The indicator may capture more than just technology adoption; for instance, it might also reflect the retraction of the most polluting process in the production function, or, in the case of multi product firms, the destruction of production of products that omit disproportionately relative to their marginal value added. Moreover, it may pick up other external shocks. To assess whether the indicator truly reflects technology adoption, I conduct various counterfactual analyses.

First, I test whether the adoption shock coincides with changes in other input variables, such as capital and labor, to ensure that the observed effect is truly linked to the adoption of new inputs. To verify that the effect does not merely reflect a general contraction in input use, I also examine whether the amount of freely allocated emission certificates remains stable during the adoption period. These allocations depend on the firmâs expected output based on its ex-ante production profile.

To assess the relationship between the adoption decision and each input variable, I employ two complementary methods. The first is an event-study design, which examines the structural dynamics of the

[27] This transitory state is often times used as moment condition in models on financial frictions and adjustment cost (Cooper and Haltiwanger, 2006; Clementi and Palazzo, 2016). I discretize this decision and add a structural state to it as a second moment condition.

indicator but requires a restricted sample. The second is a first-difference design, which uses the full sample but focuses only on short-term, transitory effects. If the freely allocated certificates remain unaffected while capital and labor respond positively, this would indicate that the observed shock is indeed technology-specific.

Second, I investigate whether the aggregate reductions identified in the reduced-form analysis - where more productive and more polluting firms reduce their emissions more - are aligned with the indicator. In other words, whether these are the same firms. I replicate the local projection regressions from the reduced-form analysis, but now use the binary adoption decision as the outcome variable. This allows me to test whether the adoption decision yields similar first-order effects as those observed in the reduced-form setting.

Finally, I explore whether carbon policy has influenced these adoption decisions. Specifically, I test whether changes in the carbon price affect the likelihood of firms adopting new technologies, thereby evaluating whether the policy serves as a meaningful driver of technology adoption.

Co-movement of other financial variables I set up a first event-study type regression to assess the relationship between key variables and the indicator of structural increases at the clean technology margin. While this analysis does not aim to identify the causal effect of clean technology shifts, I adopt the notation and framework commonly used in the difference-in-differences literature (Callaway et al., 2021).

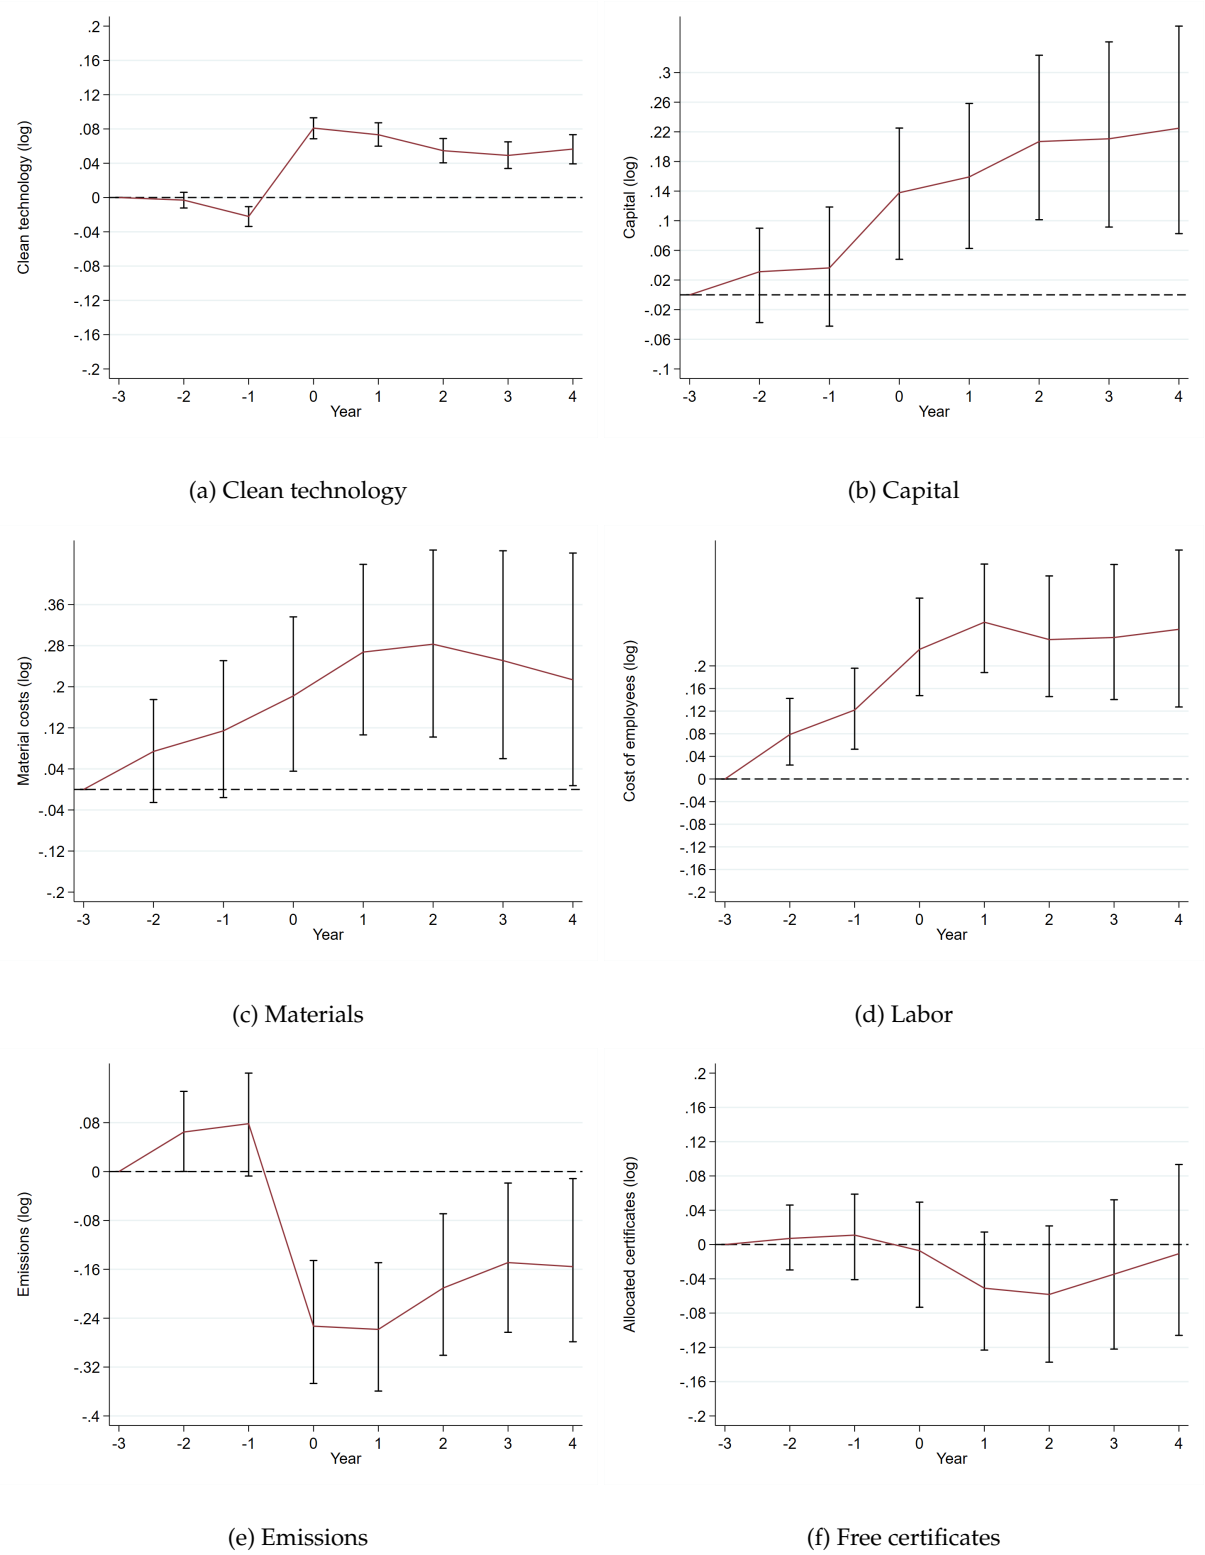
In this setting, clean technology shifts occur at staggered times, as firms choose when to adopt novel technologies. Firms engage in the decision to adopt novel technologies at time $\tau = -1$ with the outcomes at $\tau = 0$. I compare the firms that adopt increased technologies once $D_i = 1$ with those that never did so $D_i = 0$, excluding firms with multiple reductions to maintain clarity. The regression specification can be written as:

$$ln(x_{i,t}) = \sum_{\tau=-3}^{T=4} \alpha_\tau G_\tau D_i + \iota_t + \delta_i + \varepsilon_{i,t}, \tag{D.48}$$

where $x_{i,t}$ is the variable of interest - clean technology, capital, labor, amount of freely allocated certificates and emission output. The variable G_τ denotes the distance - in years - from the actual treatment decision. ι_t and δ_i denote time- and firm fixed effects. The strictest specification has time-sector-country fixed effects and captures all variation unique to the growth of a sector or country. Error terms $\varepsilon_{i,t}$ are clustered at the firm level. I approach the selection of the control group in similar spirit to Young (2019). I select a random sample of firms out of all firms that never reduce their clean technology: $\sum_t^T \omega_{i,t} = 0$. They obtain a placebo treatment $G_\tau = 1$ at a certain time τ. This procedure is bootstrapped 1000 times. The point-estimate in the graphs in Figure D.7 show the median value and the confidence intervals the $10-90$ percentile. The range of the eventstudy $t \in [-3,4]$ is chose to be exact 1 year longer in both dimensions than the structural state of the technology adoption (Equation D.47) to capture potential returns to the mean.

Figure D.7 illustrates the joint evolution of the input variables, capital, material costs, labor costs, and emissions, together with the technology adoption indicator. It also depicts the evolution of the underlying clean technology variable and the number of free certificates. The results are normalized at $t = -3$ to capture the pre-adoption state of the indicator, as defined in Equation D.47.

Figure D.7: Co-movement variables during 'technology adoption' with $\chi = 1.10$



(a) Clean technology

(b) Capital

(c) Materials

(d) Labor

(e) Emissions

(f) Free certificates

Note: Event-study type regressions for clean technology (a), capital (b), materials (c), labor cost (d), emissions (e) and the freely allocated certificates (f) following Equation D.48. The adoption occurs at time t=0.

Notably, firms that adopted emission reduction technologies exhibit an approximate 10% increase in capital, labor, and material inputs compared to firms that did not. This increase is therefore not driven by a substitution away from more expensive or emission-intensive inputs, but rather by factors that expand overall input utilization. Intuitively, when firms require new inputs, they tend to acquire the cleaner vintages of these inputs. This mechanism can be interpreted through the lens of a putty-clay model - albeit not linked to either input factor specifically - although it is not explicitly modeled here (Atkeson and Kehoe, 1999). The evolution of material costs is more complex to interpret, as materials may act either as substitutes for or complements to emission outputs.

Secondly, the emission output had not experienced any upward movement after the 2nd year post reduction - the year the constraints by the structural state of the reduction indicator had ended (i.e. W in Equation D.47). This indicates that there is no return to the mean afterwards for the reduction in emissions, providing further evidence of the indicator capturing structural change within the firm.

First-differences One major caveat of the event-study regression discussed above is that it includes only firms that adopt clean technology once, which are then compared to those that never do. To augment the relationship between adoption and changes in other input factors and include also those firms that adopt more than once, I extend the analysis using a first-difference design. Formally:

$$\Delta x_t = d_{i,t} + \iota_t + \sigma_i + \varepsilon_{i,t} \tag{D.49}$$

where $\Delta x_t = \ln(x_t) - \ln(x_{t-1})$ denotes the differences for x, defined as capital (1), emissions (2), labor cost (3), materials (4) and the number of free allocated certificates (5). The fixed effects ι_t and σ_i capture time and firm fixed effects. This approach captures whether adoption coincides with an above-average increase in inputs for all firms in the sample. However, it focuses solely on the transitory component of the emission-reduction indicator and does not account for its structural component as defined in Equation D.47.

Table D.9 provides the results for the first differences regression. The results are line with the event-study type design expressed above. A structural increase of 10% is associated with an increase in capital and labor of 5% and 6% respectively. In contrast, firm emissions decline with on average 26% [28]. Other variables such as materials and the number of freely allocated certificates seem not to be affected by the adoption.

The signaling function of price To see whether firms internalize the imposed carbon policy in their decision to adopt novel technology rather than other domestic carbon policies or other factors such as structural change and other trade shocks, I conduct a second regression analysis that examines the influence of the carbon price ζ_t on the decision to reduce for more polluting firms. Formally, the regression can be stated as:

$$d_{i,t} = \beta \ln(a_{i,t}) + \phi \ln(a_{i,t}) \ln(\zeta_t) + \iota_{c,t} + \tau_{s,t} + \delta_i + \varepsilon_{i,t} \tag{D.50}$$

[28] Note that this provides suggestive evidence that relation between emissions and capital is highly non-monotonic. Firms with higher capital levels are associated with higher levels of emissions but an increase in capital might also decrease firm level emissions

Table D.9: Relation input factors and adoption

	(1)	(2)	(3)	(4)	(5)
Variables	Δ_k	Δ_e	Δ_l	Δ_m	Δ_κ
Adoption	0.052***	-0.265***	0.063***	0.016	-0.016
	(0.015)	(0.021)	(0.013)	(0.020)	(0.010)
Observations	40,000	32,865	35,893	31,276	32,479
Adjusted R-squared	0.070	0.100	0.093	0.066	0.128
Firm fe	✓	✓	✓	✓	✓
Country-year fe	✓	✓	✓	✓	✓
Clustered S.E.	Firm	Firm	Firm	Firm	Firm

Note: Relation between the adoption decision and input variables for the whole sample of firms that participate in the EU ETS. First differences in the input variables are set using the following formula: $\Delta_x = \ln(x_t) - \ln(x_{t-1})$. The variables are capital (1), emissions (2), labor cost (3), materials (4) and the number of free allocated certificates (5). Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

where the yearly reduction decision $d_{i,t}$ depends on clean technology $\ln(a_{i,t})$ interacted with the carbon price ζ_t. clean technology is normalized to facilitate readability using a min-max approach. To account for additional variation, I include sector-year $\tau_{s,t}$, country-year $\iota_{c,t}$, and firm fixed effects δ_i. The first 2 capture the aggregate variation through rising stringency of the cap as well as potential changes in domestic policies or changes in sectoral demand. The latter fixed effect δ_i captures all firm-specific heterogeneity such that the only variation left comes from the interaction between clean technology $a_{i,t}$ and the price of carbon.

The findings in Table D.10 suggest that firms primarily consider the price of carbon certificates in their emission reduction decision, revealing the importance of ETS as a driving mechanism. While other factors like structural changes might also play a role, the carbon policy does affect the decision to invest and adopt novel technologies. The elasticity of a reduction should be interpreted as conditional to clean technology $\varepsilon = \phi * \ln(a_{i,t})$.

Ex-ante firm characteristics The mechanism in Section 2 in the main body of the paper presumes that firms that reduce their emissions are typically larger and more polluting. Reduced form evidence all hinted towards this being true. I assess whether this relation is also true for their decision to invest. I compare their ex-ante firm characteristics akin to the local projection regression in Figure B.2: one linking reductions in Phase II to their ex-ante conditions in Phase I 2007, and another for Phase III based on their 2012 characteristics, using a linear probability model:

$$d_{\bar{t}} = \beta \ln va_{i,0} + \eta \ln a_{i,0} + \iota_s + \delta_c + \varepsilon_{i,0} \tag{D.51}$$

where $d_t \in \{0,1\}$ indicates if a firm invests at least once in years $t \in \{2008\text{-}2012, 2013\text{-}2018\}$. $va_{i,0}$ and $a_{i,0}$ represent the initial value-added and clean technology, with β and η capturing their effects. Errors are clustered by sector-country, with fixed effects δ_c and ι_s accounting for cross-country policy and cross-sectoral

Table D.10: Signaling function of price

	(1)	(2)	(3)	(4)	(5)
VARIABLES	$\lambda = 10\%$	$\lambda = 10\%$	$\lambda = 10\%$	$\lambda = 5\%$	$\lambda = 20\%$
$\ln(a_i)ln(\zeta)$	-0.014**	-0.016**	-0.015**	-0.020**	-0.008
	(0.007)	(0.007)	(0.007)	(0.009)	(0.005)
$\ln(a_i)$	-0.797***	-1.367***	-0.827***	-0.992***	-0.560***
	(0.048)	(0.197)	(0.050)	(0.057)	(0.038)
$\ln(a_i)ln(P^{oil})$		0.135***			
		(0.045)			
$\ln(\kappa) - ln(\zeta)$			0.001		
			(0.002)		
Observations	18,911	18,911	18,653	18,911	18,911
Adjusted R-squared	0.081	0.082	0.083	0.071	0.068
Year-Country fe	✓	✓	✓	✓	✓
Year-Sector fe	✓	✓	✓	✓	✓
Firm fe	✓	✓	✓	✓	✓
Clustered S.E.	firm	firm	firm	firm	firm

Note: Relation between the price of carbon certificates and the decision to adopt novel levels of clean technology for more polluting firms. (1) - (3) show the specifications for a 10% increase in clean technology. (1) shows the baseline regression; (2) includes the interaction with an alternate price indicator, the global price of oil; and (3) includes the ratio of fixed certificates to total emissions as a test of the rationality of the firm. Specification (4) and (5) resume the baseline specification but for a reduction of 20% and 30% respectively. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

differences.[29]

D.2 Mapping data to the model

the logarithms of the variables are normalized using min-max scaling to preserve its non-parametric structure. To combat outliers in the model, the data is winsorized at the 1% level to reduce the impact of outliers. The scaling is the following:

$$x^{norm} = \frac{x - x^{min}}{x^{max} - x^{min}} \tag{D.52}$$

Where x is the logarithm of the value if the value is not yet expressed as logarithm. x^{norm} is therefore the normalized variable. Except for the innovation of the transitory component z_{it}^T which is i.i.d. and has a normal distribution $N(0, \sigma_z^2)$, the distributions for the underlying state variables are non-parametric, as is their joint distribution.

The data for the 3 underlying state variables is binned into categories where the size of the grid for the 3 values $a^{grid} = 306$ and $z^{F,grid} = 5$. The grids are spaced in equal distances of the logarithm of the variable (note that z^T and z^F are already expressed in logs). Each firm is assigned a grid value based on the minimal

[29] Country fixed effects are crucial due to variation in carbon policies. For example, Sweden complements the ETS with national measures, while Belgium offsets relocation risks with subsidies.

Table D.11: Ex ante characteristics adopters (Phase II)

	(1)	(2)	(3)	(4)	(5)
VARIABLES	$\lambda = 10\%$	$\lambda = 5\%$	$\lambda = 20\%$	$\lambda = 10\%$	$\lambda = 10\%$
nrl_prf	0.053	0.325***	0.020		
	(0.070)	(0.101)	(0.055)		
nrl_ct	-0.722***	-0.725***	-0.631***	-0.683***	-0.667***
	(0.077)	(0.091)	(0.084)	(0.075)	(0.076)
lnva				0.012*	
				(0.007)	
lnk					0.010*
					(0.006)
Observations	1,801	1,801	1,801	1,883	1,876
Adjusted R-squared	0.085	0.091	0.055	0.085	0.085
Country-Sector fe	✓	✓	✓	✓	✓
Clustered S.E.	Country-Sector	Country-Sector	Country-Sector	Country-Sector	Country-Sector

Note: Relation between the ex-ante characteristics and the cumulative probability to adopt during phase II of the adoption. The ex-ante characteristics are base on 2007. The indicator is 1 if the firm adopts at least one during the period from 2008-2012. (1) - (3) show the specifications for different threshold values of λ. Specification (4) and (5) resume the baseline specification but use capital and value added as proxy variables for size. Clustered standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table D.12: Ex ante characteristics adopters (Phase III)

	(1)	(2)	(3)	(4)	(5)
VARIABLES	$\lambda = 10\%$	$\lambda = 5\%$	$\lambda = 20\%$	$\lambda = 10\%$	$\lambda = 10\%$
nrl_prf	0.268***	0.515***	0.233***		
	(0.100)	(0.101)	(0.085)		
nrl_ct	-0.886***	-0.886***	-0.705***	-0.881***	-0.807***
	(0.098)	(0.110)	(0.085)	(0.096)	(0.087)
lnva				0.042***	
				(0.012)	
lnk					0.033***
					(0.010)
Observations	1,843	1,843	1,843	1,918	1,912
Adjusted R-squared	0.102	0.133	0.097	0.106	0.102
Country-Sector F.E.	✓	✓	✓	✓	✓
Clustered S.E.	Country-Sector	Country-Sector	Country-Sector	Country-Sector	Country-Sector

Note: Relation between the ex-ante characteristics and the cumulative probability to adopt during phase II of the adoption. The ex-ante characteristics are base on 2012. The indicator is 1 if the firm adopts at least one during the period from 2013-2018. (1) - (3) show the specifications for different threshold values of λ. Specification (4) and (5) resume the baseline specification but use capital and value added as proxy variables for size. Clustered standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

distance between the bin and the value. This ensures the non-parametric nature of the joint distribution between a and z^F, which arises naturally. Figure D.8 documents a positive cross-sectional relationship between clean technology and productivity, confirming findings in the literature (Copeland et al., 2022; Shapiro and Walker, 2018). The relation between the state variables and other variables - either reduced form or within the context of the model - is discussed in other parts of the paper.

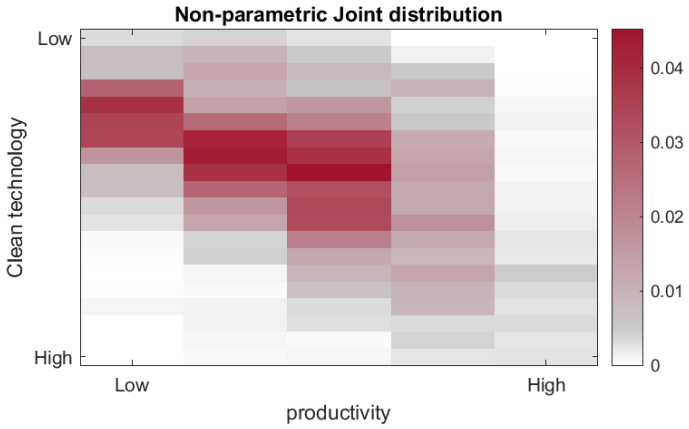


Figure D.8: Joint distribution at t=0

Note: Joint distribution of firms at the start of the transition. The distribution is split up in 20 evenly sized grids of clean technology and 5 evenly sized grids of productivity. The more red the gridblock becomes, the higher the density of firms in that grid.

D.3 Trade cost τ

Finally, the estimation of the model requires the parameterization of the trade cost τ. In this setting, I base myself on Arkolakis et al. (2012): Given certain restrictions, one can state that the ratio of exported goods relative to goods produced for the domestic market equals $\tau^{1-\eta}$. These restrictions are threefold: (i) balanced trade ; (ii) aggregate profits as a constant share of aggregate revenues; and (iii) a CES demand system with a constant elasticity of trade with respect to variable trade. If all these three conditions are met, the relative importance of the trade channel is determined by the iceberg trade cost τ:

$$\frac{\int_{z*a} p_x(z,a) q_x(z,a) d\mu_h(z,a)}{\int_{z*a} p_d(z,a) q_d(z,a) d\mu_h(z,a)} = \frac{\tau^{1-\eta} \int_{z*a} p_d(z,a)^{1-\eta} P_f^{\eta} I_f d\mu_h(z,a)}{\int_{z*a} p_d(z,a)^{1-\eta} P_h^{\eta} I_h d\mu_h(z,a)} \tag{D.53}$$

Prior to the introduction of the Pigouvian carbon policy, I assume that the firms behave as if they are in balanced trade. Along with two other conditions that are met throughout the model, the aggregate export share is a sufficient condition to determine the trade cost τ.

The aggregate export and production statistics are gathered from the World Input Output Tables (WIOT) for the year 2005. I aggregate all manufacturing sectors to 1 sector and all countries that are part of the EU ETS (as well as Great Britain) are aggregated into 1 domestic country relative to the Rest of the world. I use

the Head-Ries index (HRI) (Head and Ries, 2001; Antras and Chor, 2018) that obtains bilateral trade costs from the data. The measure cancels out potential differences in aggregate resistance terms and prices in the data. The trade cost can be stated as :

$$\tau = \left(\frac{X_{hx} X_{fx}}{X_{h,d} X_{h,f}} \right)^{\frac{1}{1-\eta}} \tag{D.54}$$

where X denotes the total supply of goods for each market. The demand elasticity in the model is set to $\eta = 6$. Note that throughout the model, I make the assumption that the trade cost is relative to that of the domestic country (i.e. normalized to 1).The total iceberg trade cost for 2005 amounts to $\tau = 1.26$.

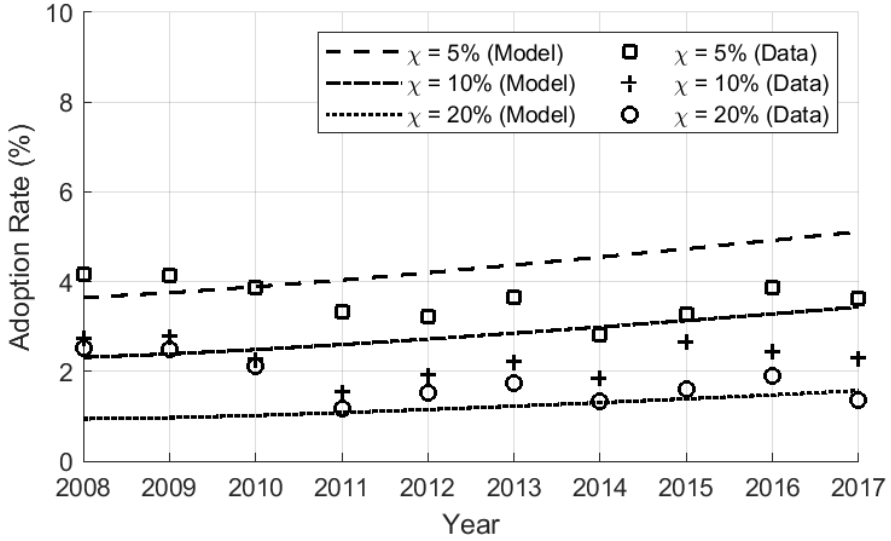
E. Results Structural Model

This section complements the results in Section 5 in the main text by presenting further results. It includes additional validation exercises for the 2008 - 2012 period (Phase II of the EU ETS) and reports outcomes for aggregate domestic revenues instead of aggregate domestic demand, along with their first-order decomposition.

E.1 Validation Exercises

This section provides additional validation of the model. I first focus on the adoption decision taken by the firm. The main paper provides the evolution of the price of carbon and the cumulative probability of adoption between the period of 2013 and 2018 - phase III of the ETS. in both the data and the model as sufficient conditions to validate the model. This section provides further evidence that the model does a good job in explaining aggregate differences. I first focus on the technology adoption indicator, showing both the evolution over time as well as the decomposition by both margins - clean technology and productivity. Then I show that the model does not stop there, but it also can explain evolutions in the outcome variables.

Figure E.9: Adoption rate (2008 - 2017)



Note: The decomposition of the evolution of aggregate domestic sales relative to the business-as-usual scenario. The dashed line shows the total effect of the carbon policy. It can be decomposed in two components. The direct effect of the carbon policy (blue) is calculated under a fixed distribution of firms. The adoption effect (red) captures changes due to shifts in the firm distribution. The dashed line is the exact sum of the two components.

Technology adoption over time The evolution of technology adoption between 2008 and 2017 is illustrated in Figure E.9. The lines represent the model's predicted outcomes, while the square, cross, and circle markers denote the observed rates of clean technology adoption derived from the data (respectively 20%,

10% and 5%). The model outcomes exhibit an upward trend which is driven by two key factors: first, the increasingly stringent aggregate conditions that raise the carbon price and thereby incentivizing firms to adopt. Second, the effect of clean technological progress, which gradually reduces the cost of adoption over time.

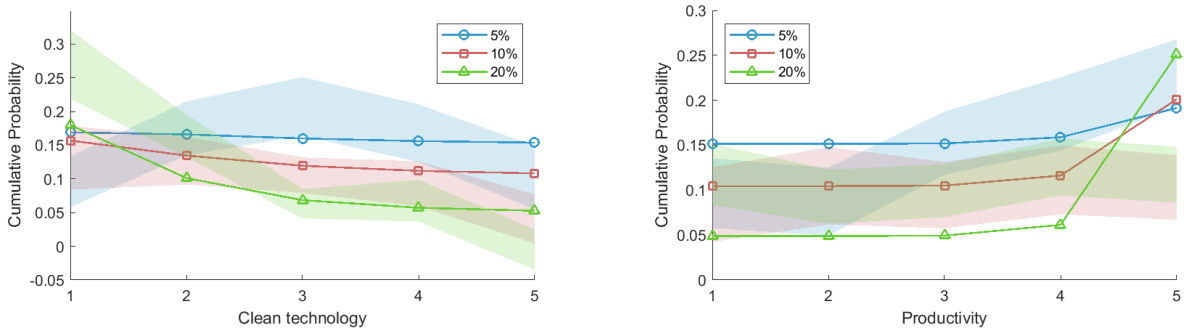
As shown in Figure E.9, the model tends to underestimate adoption rates relative to the data during the initial years of the policy. This discrepancy may stem from several factors. Potentially due to the distortionary effect of the financial crisis that overestimates increases in clean technology or through which firms had to retract some output that they would otherwise not have done. Otherwise, this could also be a case of the model underestimating the presence of a few simple idiosyncratic low-hanging fruit adoption decisions conducted by forward-looking firms. This might directly explain the slight overestimation of the adoption decision later on as a counterbalancing effect. Nonetheless, the model provides a good fit to the actual data.

E.1.1 Firm inequality

This section examines both heterogeneous implications of the carbon policy on the adoption decision as well as the evolution of emissions along the two different margins: clean technology and productivity.

Technology adoption phase II I start by examining the adoption decision for phase II of the ETS. Figure E.10 shows the cumulative probability of adoption between 2008 and 2012 along the two margins of heterogeneity - clean technology and productivity for the data (shaded area) and the model (in lines).

Figure E.10: Idiosyncratic decisions by dimension (2008 - 2012)



(a) Clean technology margin

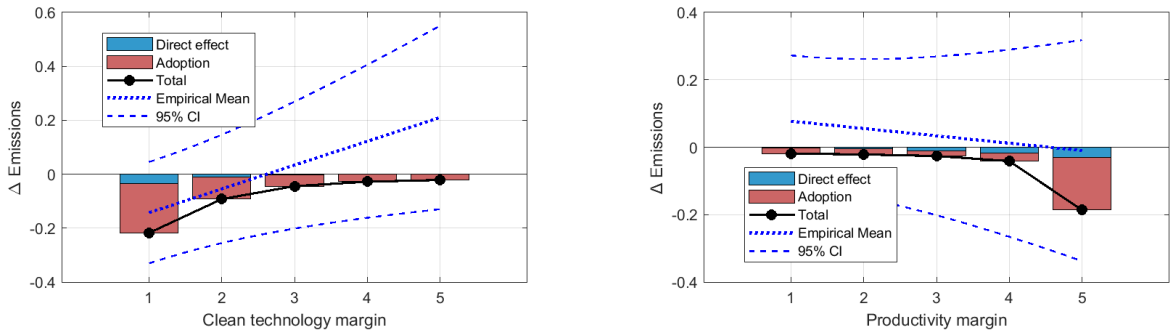
(b) Productivity margin

Note: Cumulative probability of adoption between 2008 and 2012 along 2 dimensions of heterogeneity: the clean technology margin (left panel) and the productivity margin (right panel). They are considered as a Bernoulli trail: $1-\prod_{t=2013}^{2018}(1-d_t(\chi))$. The margins are equally spaced bins based on the ex-ante levels of clean technology and productivity (i.e. at the year 2008). The more right on the x-axis the higher average clean technology / productivity of the quintile. the colors show the adoption probabilities for the 3 different values of $\chi = \{1.05, 1.10, 1.20\}$. The shaded area represents the 95% confidence intervals in the data for their respective decisions given the equivalent ex-ante value of the margin, conditional on the average level of the other margin.

The results in Figure E.10 indicate a good fit between the data and the model, though slightly weaker than in the case of phase III. As discussed previously in relation to Figure E.9, this discrepancy may stem from the effects of the financial crisis and/or unaccounted 'low-hanging fruit' in the model, reflecting additional unobserved firm-level heterogeneity not captured by the current specification. Similar to the situation in phase III, the model slightly underestimates adoption probabilities for the most polluting firms (those with clean technology 1) and overestimates them for the most productive firms (in the fifth productivity quintile).

Inequality in outcome variables Figure E.11 presents the evolution of emissions over the period 2012-2018, decomposed along two key dimensions of heterogeneity: the clean technology margin and the productivity margin. This exercise serves as an additional validation of the model's performance. The empirical evolution of emissions, shown by the blue lines and associated confidence intervals, aligns well with the model-implied patterns.

Figure E.11: Emission reduction by dimensions (2012-2018); Validation



(a) Clean technology a_t

(b) Productivity a_t

Note: Evolution of emissions reduction along 2 dimensions of heterogeneity: clean technology margin (left panel) and the productivity margin (right panel). The margins are equally spaced bins based on the ex-ante levels of clean technology and productivity (at the year 2008). The more right on the x-axis the higher average clean technology / productivity of the quintile. The reductions are split into a direct effect and an adoption component as in the baseline figure. The blue lines show the average evolution of emissions along the margins present in the data. The dashed line depicts their confidence intervals.

E.2 The effect of carbon policy on aggregate variables

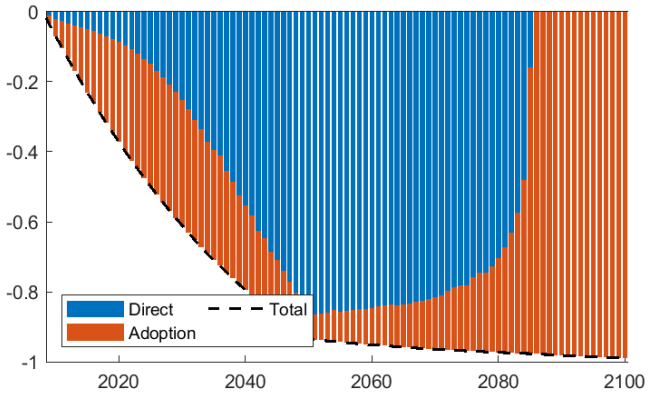
The baseline results discussed in the main text focus exclusively on the effects of the carbon policy on domestic demand, which is the primary variable of interest for a welfare-maximizing domestic government. This section broadens the analysis by examining the evolutions of other key aggregate variables to provide additional intuition. First, I decompose the evolution of aggregate domestic emissions over time. Second, I analyze the sales of domestic producers to illustrate the business-stealing effect by foreign competitors. Finally, I examine the evolution of foreign emissions, since overall welfare depends not only on domestic demand but also on emissions as a negative externality.

E.2.1 Domestic emissions

The evolution of aggregate emissions is illustrated in Figure E.12. As with aggregate demand, it can be decomposed into two channels: the direct effect of the carbon price on output (blue) and the effect of changes in the distribution (red). The decomposition follows the same approach described in Equation 22. Both channels contribute negatively to total domestic emissions, since $e = q/a$ and $E = \int e$.

The relationship between adoption and the output retraction channel is less straightforward than in the case of aggregate demand. Over time, the progressively higher carbon price required to clear the emissions market leads to a growing share of emissions reductions that would have occurred even under a fixed firm distribution. Consequently, the contribution of the direct channel increases up to the year 2050, after which it gradually decreases to 0.

Figure E.12: Decomposition of domestic emissions



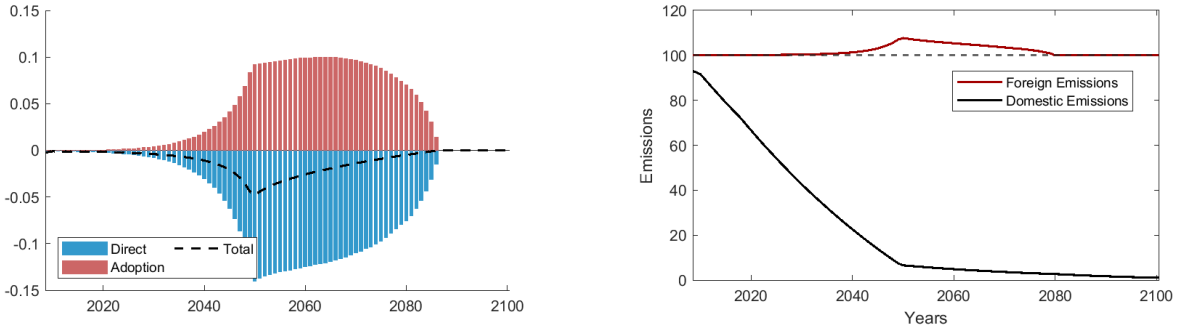
Note: decomposition of the evolution of aggregate domestic emissions relative to the business-as-usual scenario. The dashed line shows the total effect of the carbon policy, decomposed into the direct effect (blue) under a fixed firm distribution and the adoption effect (red), capturing changes due to shifts in the firm distribution.

Meanwhile, the evolution of aggregate quantities reveals that as emissions become increasingly decoupled from output, due to the mitigating effect of adoption, firms experience smaller adverse impacts on production. As time progresses, like the decomposition of demand, all reductions in emissions are driven

through the adoption channel.

E.2.2 Domestic sales

Figure E.13: Aggregate sales and foreign emissions (relative BAU)



(a) Aggregate domestic sales

(b) Foreign emissions

Note: Panel (a) shows the decomposition of the evolution of aggregate domestic sales relative to the business-as-usual scenario. The dashed line shows the total effect of the carbon policy, decomposed into the direct effect (blue) under a fixed firm distribution and the adoption effect (red) capturing changes due to shifts in the firm distribution.The dashed line is the exact sum of the two components. Panel (b) displays the corresponding decomposition for foreign emissions (red) and domestic emissions (black) relative to the BAU scenario. Note that the evolution of domestic emissions in this panel is relative to the BAU, such that it differs slightly from the carbon cap in the text.

This section examines the evolution of aggregate domestic sales. Figure E.13a presents domestic sales over time, expressed relative to a business-as-usual (BAU) scenario. The interpretation of the BAU benchmark differs slightly in this context, as both domestic and foreign incomes are normalized to one. Under perfect symmetry (i.e., in the absence of growth differentials or distortions) the sum of all foreign and domestic sales equals the sum of foreign and domestic income. Since total income is normalized to unity, deviations in domestic sales shown in the figure fully capture the business-stealing effect from foreign firms operating in the domestic market. The effects on sales are a bit more severe than those for aggregate demand as they close off the substitution channel that demand can rely on.

E.2.3 Foreign emissions

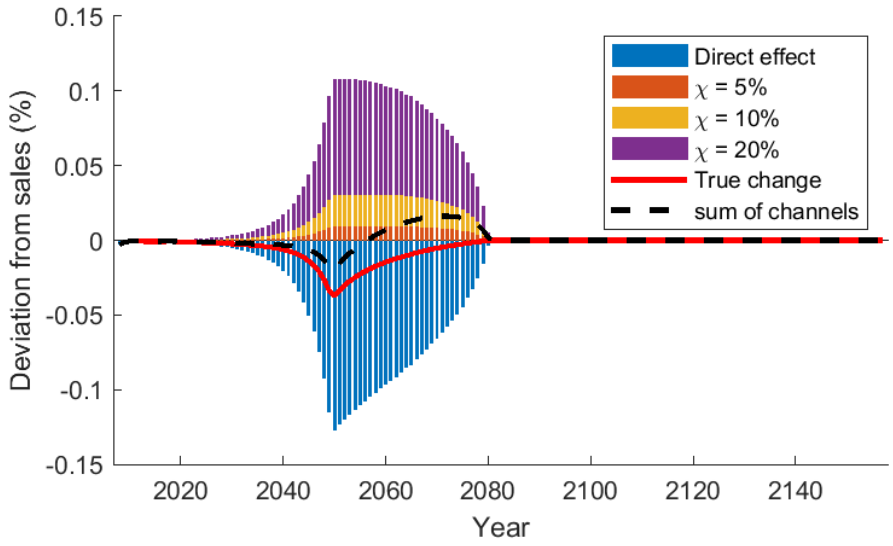
The evolution of foreign emissions relative to the BAU scenario is shown in Figure E.13b. While domestic emissions decline to around 10% of their initial level by 2050 and gradually approach zero thereafter, foreign emissions increase substantially, peaking at approximately 9% in 2050. This rise is entirely driven (by construction) through the reallocation effect induced by changes in general equilibrium prices.

E.3 Decomposing up to a first-order approximation

This section further decomposes the aggregate output and aggregate sales variables discussed in the main text. It separates the adoption channel, illustrating the relative contributions of large (20%), medium (10%), and minor (5%) adoption decisions, and the output channel, highlighting the role of the foreign market in firms' decisions to reduce output. The corresponding mathematical derivations are provided in Appendix C.

As the decompositions are only an approximation up to the first order of the aggregate price indices, they are quite far off the true value when the impact of the policy is substantial. It is therefore only recommended to use these as indicative to have a sense of the relative magnitude rather than to interpret that as absolute numbers.

Figure E.14: Decomposition of the adoption channel



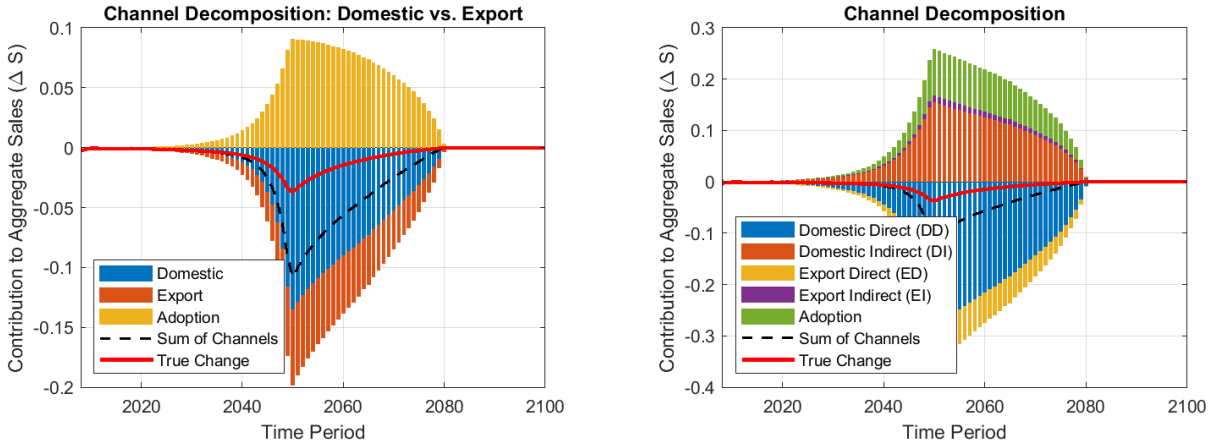
Note: decomposition of the evolution of aggregate domestic sales relative to the business-as-usual scenario. The dashed line shows the total effect of the carbon policy. It can be decomposed in two components. The direct effect of the carbon policy (blue) is calculated under a fixed distribution of firms and the adoption effect, which captures changes due to shifts in the firm distribution. The latter is further decomposed up to a first-order approximation in either of the 3 decisions of 5%, 10% or 20%.

Decomposition of the adoption margin The adoption channel can be further decomposed into the individual contributions of the three key decisions, which account for approximately 20%, 10%, and 5% of the overall effect (see Figure E.14). It is worth noting that the total estimated change does not exactly match the true change, primarily because the decomposition relies on a first-order approximation and the underlying variations in ΔS are relatively large. In relative terms, the results suggest that the overall impact is driven predominantly by firms adopting the higher probability levels of 20% rather than the more incremental adjustments of 5% (the latter being more common in the data). One reason is the fact that the more polluting

and productive firms are those that are more inclined to adjust with 20% as for them the payoff is higher.

Decomposition of the direct margin The output channel can be further decomposed to distinguish the relative contributions of the trade and domestic channels. The domestic channel is substantially larger than the trade channel, indicating that the latter plays a limited role in determining the overall implications for aggregate losses.

Figure E.15: Decomposition of domestic sales



(a) Decomposition trade channel (i)

(b) Decomposition trade channel (ii)

Note: Decomposition of aggregate sales into a component for goods catered to the domestic market and goods catered towards the domestic market. The left panel (i) shows the decomposition of the price component of aggregate sales into the share meant for export and the share for domestic demand. The sum of the first order differences are shown in dashed black while the true change is given in red. The right panel (ii) decomposes aggregate sales further into the price effect (indirect) and the output effect (direct) for either component.

F. Sensitivity analysis and Extensions

This section provides additional documentation on the assumptions surrounding the extensions in Section 5. In the paper, I elaborate the 3 extensions discussed in the text: (1) one where clean technological progress depends on demand; (2) where there are 2 industries which are different in their exposure to trade and have differently determined endogenous costs of adoption; and (3) where there is aggregate uncertainty to mimic the effect of the financial crisis in 2008.

F.1 Sensitivity analysis

F.1.1 The ETS endgame

There consists substantial uncertainty on how the ETS endgame will look like exactly. For instance, net zero in the aggregate does not necessarily imply net zero in this sector as well. Moreover, uncertainty surrounding the interaction between banked certificates, financial frictions, and coordination failures could generate persistent deviations from the cap trajectory. Likewise, technological bottlenecks and political compromises may delay or reshape the effective stringency of the cap in later years.

In the baseline model in the text, I solve this bluntly and simply assume that at net zero the cap at time T is at 10% of its value at time t_0 (i.e. $E_{h,T} = 0.1E_{h,0}$). From time T onward, emissions decline at a yearly rate of 1% until full decarbonization is achieved. To test how sensitive the results are to the particular calibration of the endgame, I redo the model for three different cutoff values of $E_{h,T}$ as a function $E_{h,0}$: 5%, 10%, and 25%. The post-T trajectory of emissions remains unchanged across these variations.

F.1.2 Persistence productivity shocks

The baseline result yields the additional restriction that there is no persistence in the transitionary component $z^T_{i,t} \sim N(0, \sigma_z)$ to reduce the number of state variables to be tracked over time. I augment the framework and allow for the transitory component of productivity to move according to ann AR(1) process. The idiosyncratic productivity process can be stated as.

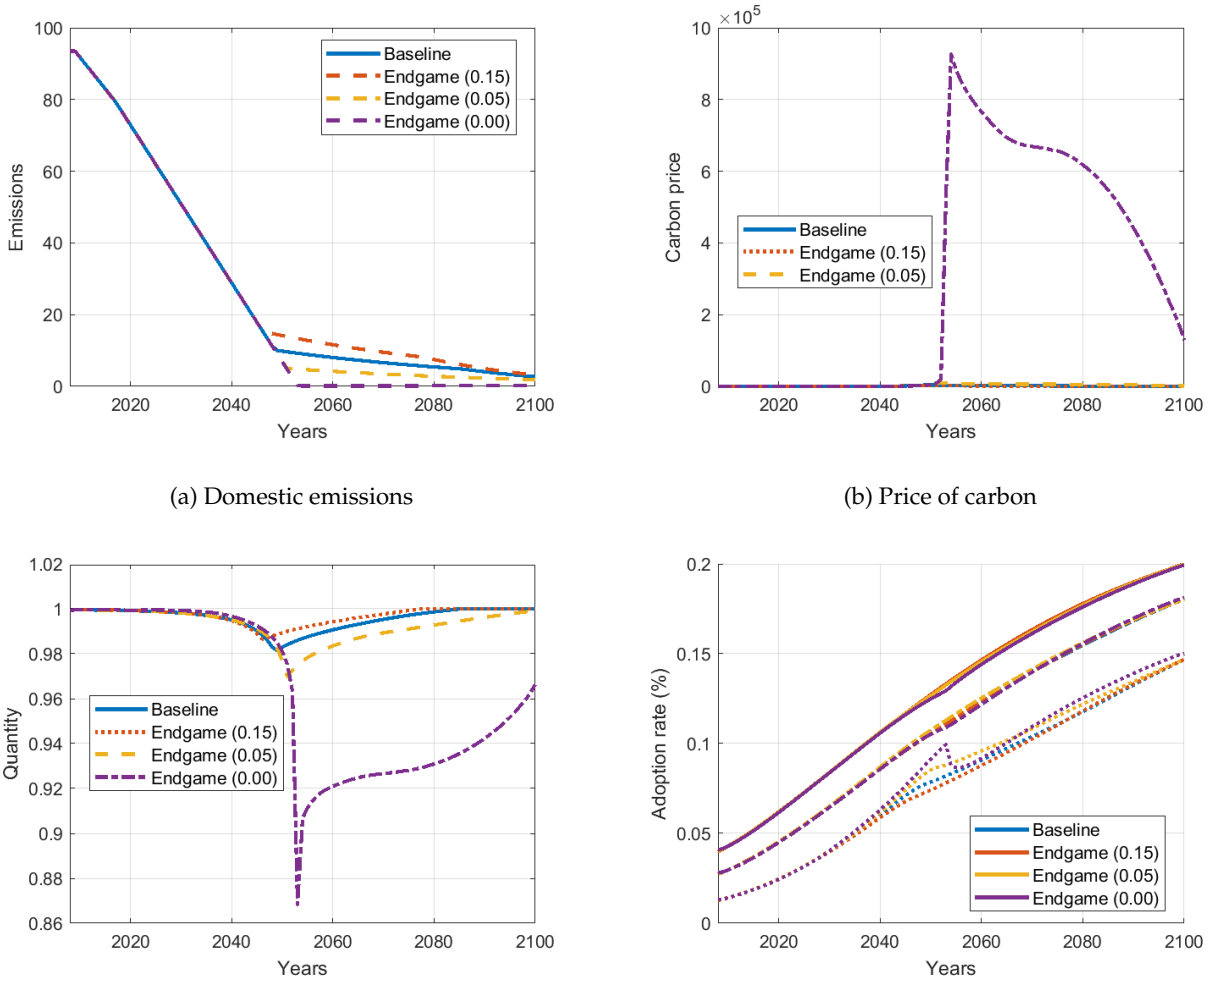
$$z_{i,t} = z^F_i + z^T_{i,t}, \qquad \text{with} \quad z^T_{i,t} = \rho z^T_{i,t-1} + \xi_{i,t} \tag{F.55}$$

where $|\rho| < 1$ and the innovation to the transitory component is i.i.d. and behaves as $\xi_{i,t} \sim N(0, \sigma^2_\xi)$. Substituting the second part of the expression for $z^T_{i,t}$, I can rewrite productivity as:

$$\begin{aligned} z_{i,t} &= z^F_i + \rho(z_{i,t-1} - z^F_i) + \xi_{it} \\ z_{i,t} &= (1-\rho)z^F_i + \rho z_{i,t-1} + \xi_{i,t} \end{aligned} \tag{F.56}$$

consequently, this equation equals the earlier Markov process as the functional form can be written as $g(z_{i,t-1}) = (1-\rho)z^F_i + \rho z_{i,t-1}$. Note that this requires a revaluation of the fixed component as well as it now is a function of the transitory component $(1-\rho)z^F_i$. The variance of the transitory component yields the

Figure F.16: Implications of endvalue $E_{h,T}$



(a) Domestic emissions

(b) Price of carbon

(c) Domestic demand

(d) Yearly Adoption rate (%)

Note: Evolution of the outcome variables for different endgame strategies: $E_{h,T} = \{0, 5\%, 10\%, 20\%)E_{h,0}$ in purple, yellow, blue and red respectively. The blue lines represent the outcome values for the baseline model used in the main body of the paper. The upper panels present aggregate emissions and the corresponding price of carbon. The lower two panels present the loss in domestic demand and the yearly adoption rate for all 4 scenarios.

common:

$$\text{Var}(z_t) = \frac{\sigma_\xi^2}{1-\rho^2}.$$

Estimation the unknown persistence parameter ρ in Equation F.56 can be directly estimated from the data:

$$\omega_{i,t} = \rho\,\omega_{i,t-1} + \psi_i + \iota_t + \varepsilon_{i,t} \tag{F.57}$$

where ψ_i and ι_t capture the fixed firm-level component and general equilibrium fixed effects; and the error terms $\varepsilon \sim N(0, \sigma_\varepsilon)$. Note that this equation is formulated using $\omega_{i,t}$ instead of $z_{i,t}$ due to the absence of price data. If the induced error margin of the price is zero, both components ω and ε would equate $z_{i,t}$ and ξ exactly. Provided the error term, the regression results are shown in *Table F.*13:

Table F.13: Simulated method of moments

Parameter	ρ	Observations	σ_ε
Value	0.41	27.923	0.32

F.1.3 Other

Finally, I assess the robustness of the results to alternative values of several externally specified (or estimated) parameters: (i) the iceberg trade cost τ, (ii) the foreign wage W_f, (iii) the discount factor β, and (iv) the elasticity of substitution between goods, η.

For all of these four parameters, I compare the baseline model with a different version in which the parameter is either increased with a factor of 5% or decreased with the same factor - except for τ for which i choose a factor of 20% to explicitly demonstrate differences. All other parameters - also the estimated ones - remain constant. The results are summarized in the bullet points below:

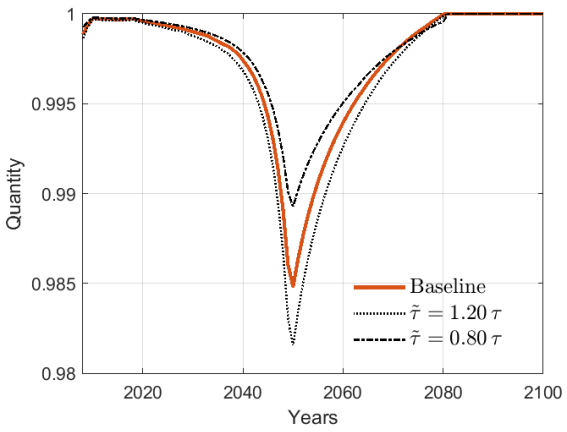
- Symmetric trade cost (τ): Changes in τ have only minor effects on the adoption decision due to the gradual implementation of the carbon policy. They therefore leave the relation between the carbon policy and overall domestic demand largely unchanged. This relation is consistent with limited reallocation effects found in the data (see Figure F.17).

- Foreign wage (W_f): Variations in W_f have, similarly to the trade cost τ little impact on the adoption decision and therefore limited impact on the relation between the carbon policy and aggregate prices/output. The exact results are shown in Figure F.18

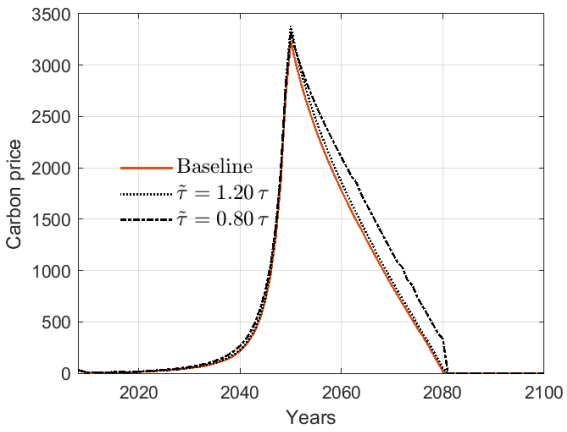
- Discount factor (β): A higher β increases the adoption rate and reduces output losses, as agents value future gains more and invest in adoption earlier (Exact point estimates for each year are shown in Figure F.19).

- Elasticity of substitution (η): A higher η slightly raises the adoption rate but amplifies output losses, since greater substitutability intensifies competition and strengthens price effects (see Figure F.20).

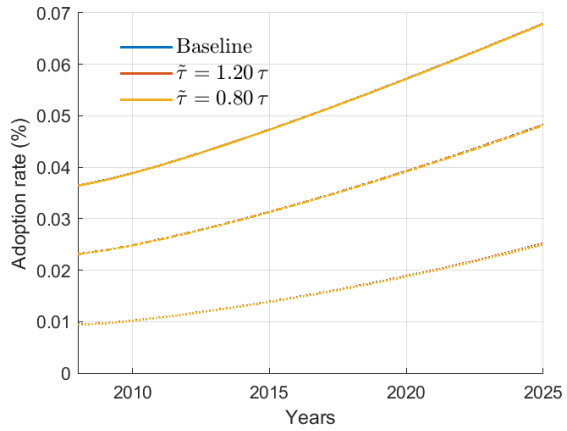
Figure F.17: Implications of iceberg trade cost τ



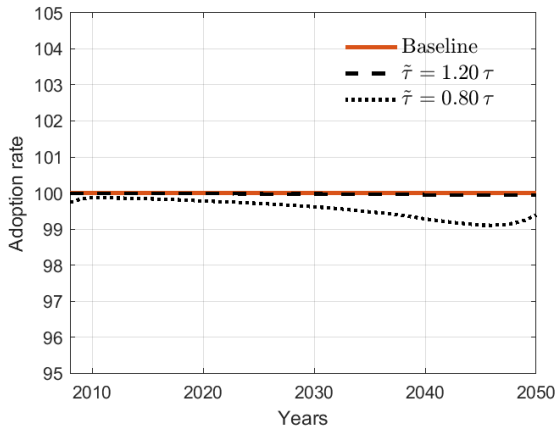
(a) Domestic demand



(b) Price of carbon



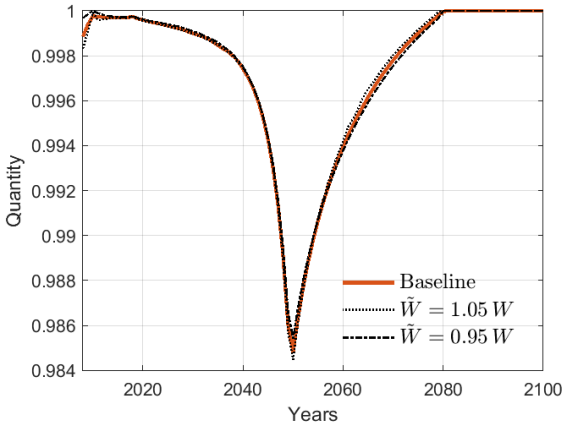
(c) Yearly Adoption rate (%)



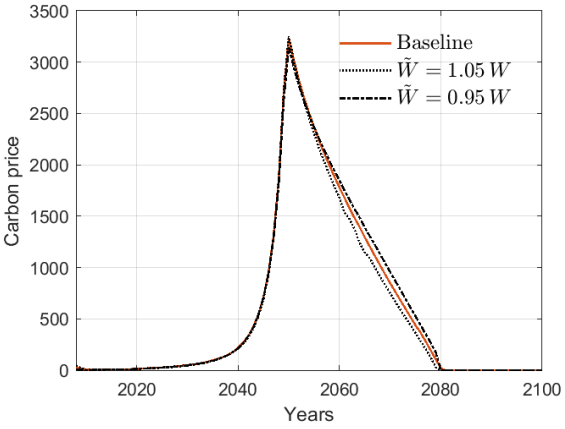
(d) Yearly Adoption rate (%)

Note: Evolution of the outcome variables given sensitivity analysis for $\tilde{\tau} = \{0.95, 1, 1.05\}\tau$. Panel (a) and (b) show the aggregate outcome variables of demand and the price of carbon respectively for the baseline (red) the 5% increase (dotted) and the 5% decrease (dashed). Panel (c) and (d) show the changes in adoption rate for either. the full line represents a $\chi = 1.20$, the dashed shows $\chi = 1.10$ and dotted $\chi = 1.05$.

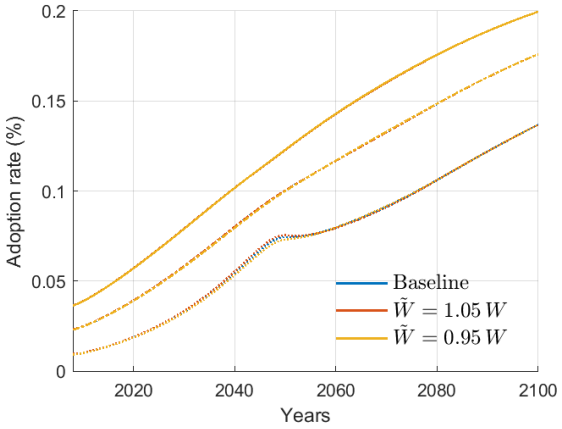
Figure F.18: Implications of foreign wage W_f



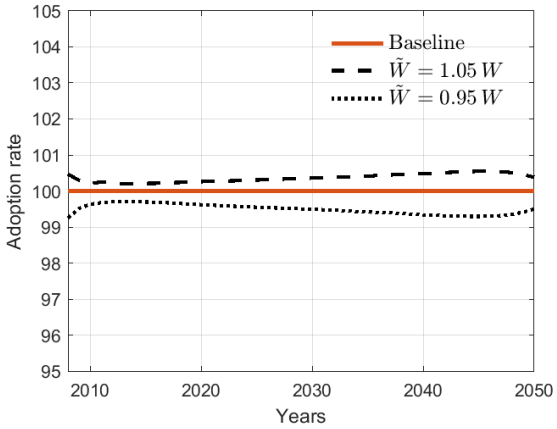
(a) Domestic demand



(b) Price of carbon



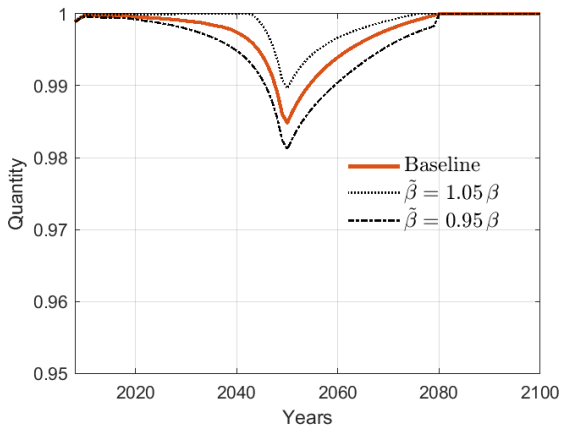
(c) Yearly Adoption rate (%)



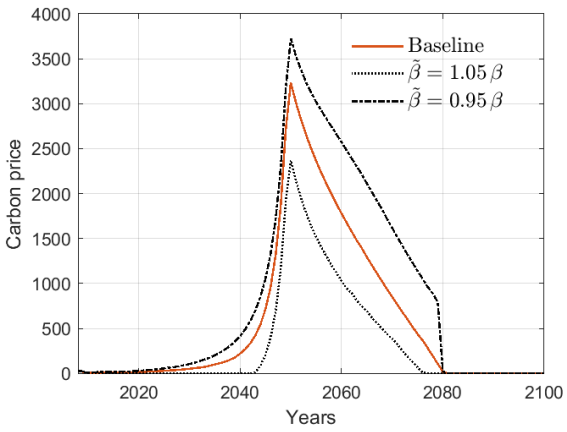
(d) Yearly Adoption rate (%)

Note: Evolution of the outcome variables given sensitivity analysis for $\tilde{W} = \{0.95, 1, 1.05\}W$. Panel (a) and (b) show the aggregate outcome variables of demand and the price of carbon respectively for the baseline (red) the 5% increase (dotted) and the 5% decrease (dashed). Panel (c) and (d) show the changes in adoption rate for either. the full line represents a $\chi = 1.20$, the dashed shows $\chi = 1.10$ and dotted $\chi = 1.05$.

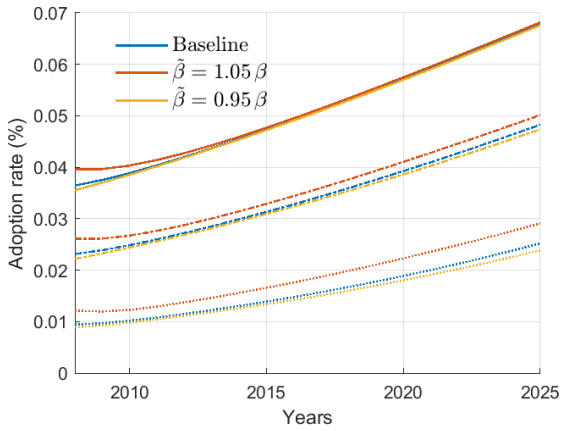
Figure F.19: Implications of discount factor β



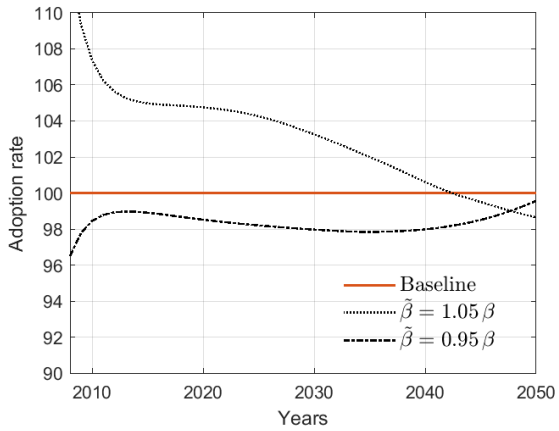
(a) Domestic demand



(b) Price of carbon



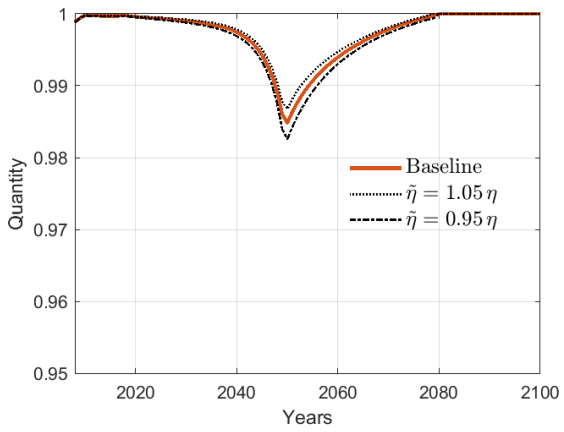
(c) Yearly Adoption rate (%)



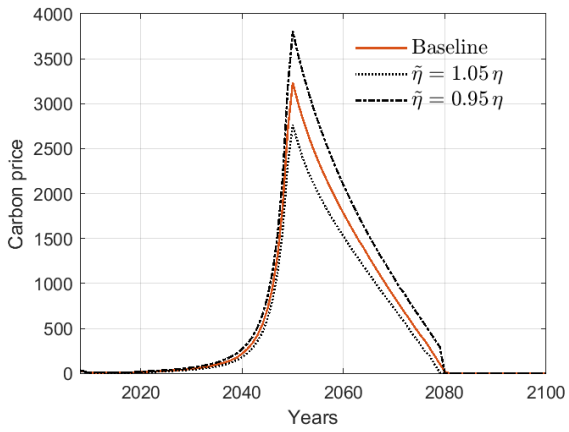
(d) Yearly Adoption rate (%)

Note: Evolution of the outcome variables given sensitivity analysis for $\tilde{\beta} = \{0.95, 1, 1.05\}\beta$. Panel (a) and (b) show the aggregate outcome variables of demand and the price of carbon respectively for the baseline (red) the 5% increase (dotted) and the 5% decrease (dashed). Panel (c) and (d) show the changes in adoption rate for either. the full line represents a $\chi = 1.20$, the dashed shows $\chi = 1.10$ and dotted $\chi = 1.05$.

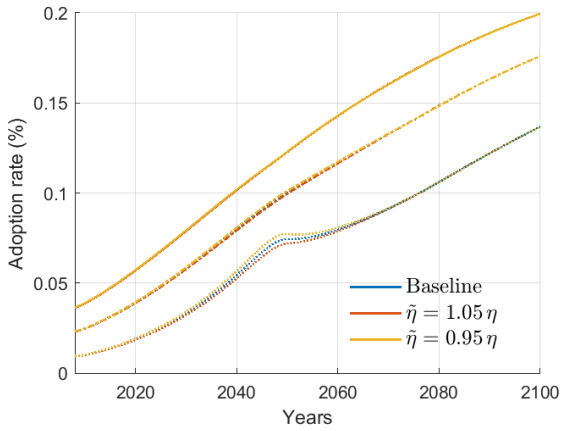
Figure F.20: Implications of substitution parameter η



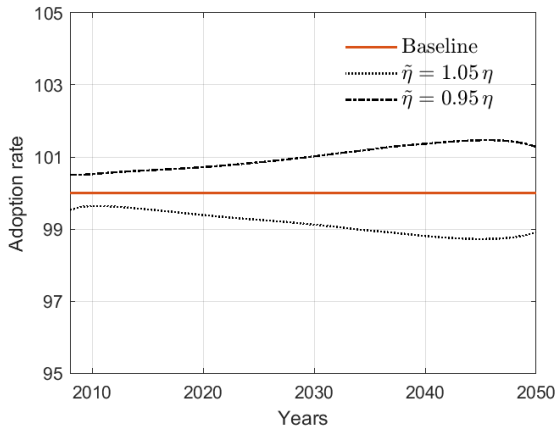
(a) Domestic demand



(b) Price of carbon



(c) Yearly Adoption rate (%)



(d) Yearly Adoption rate (%)

Note: Evolution of the outcome variables given sensitivity analysis for $\tilde{\eta} = \{0.95, 1, 1.05\}\eta$. Panel (a) and (b) show the aggregate outcome variables of demand and the price of carbon respectively for the baseline (red) the 5% increase (dotted) and the 5% decrease (dashed). Panel (c) and (d) show the changes in adoption rate for either. the full line represents a $\chi = 1.20$, the dashed shows $\chi = 1.10$ and dotted $\chi = 1.05$.

F.2 Policy instruments

This section provides additional information on the policy instruments in the text and gives economic intuition on their interaction with aggregate state variables and the idiosyncratic adoption decision. There are 2 specific policies considered. First I develop a Carbon border adjustment mechanism (CBAM) where imported goods are affected exactly the same way as the goods produced in the domestic market. Second, I let the gains from the carbon policy rebate as adoption subsidies rather than as free fixed certificates κ. The notes in this section provide further documentation to the main paper.

F.2.1 Carbon Border adjustment mechanism

Holding all other variables fixed; an introduction of the CBAM would negatively affect domestic demand even further. Substitution from these firms to the

$$\begin{aligned} P_h &= \left[\left(\frac{\eta}{\eta-1} W \right)^{1-\eta} \int_{z*a} \left(\frac{1}{z} + \frac{\zeta}{a} \right)^{1-\eta} d\mu_h(z,a) + \left(\frac{\eta}{\eta-1} W \right)^{1-\eta} \tau^{1-\eta} \int_{z*a} \left(\frac{1}{z} + \frac{\zeta^{CBAM}}{a} \right)^{1-\eta} d\mu_f(z,a) \right]^{\frac{1}{1-\eta}} \\ &= \frac{\eta}{\eta-1} W \left[\int \left(\frac{1}{z} + \frac{\zeta}{a} \right)^{1-\eta} d\mu_h + \tau^{1-\eta} \int \left(\frac{1}{z} + \frac{\zeta^{CBAM}}{a} \right)^{1-\eta} d\mu_f \right]^{\frac{1}{1-\eta}} \end{aligned} \tag{F.58}$$

Under the assumption that the firms in the foreign country are exposed to exactly the same policy as those in the domestic country, then $\zeta = \zeta^{CBAM}$. One can take the derivative to ζ as:

$$\frac{\partial P_h}{\partial \zeta} = P_h \frac{\int \frac{1}{a} \left(\frac{1}{z} + \frac{\zeta}{a} \right)^{-\eta} d\mu_h + \tau^{1-\eta} \int \frac{1}{a} \left(\frac{1}{z} + \frac{\zeta}{a} \right)^{-\eta} d\mu_f}{\int \left(\frac{1}{z} + \frac{\zeta}{a} \right)^{1-\eta} d\mu_h + \tau^{1-\eta} \int \left(\frac{1}{z} + \frac{\zeta}{a} \right)^{1-\eta} d\mu_f} > 0$$

If not, the derivative of P_h w.r.t. ζ^{CBAM} yields the following:

$$\frac{\partial P_h}{\partial \zeta^{CBAM}} = \left(\frac{\eta}{\eta-1} \tau W \right)^{1-\eta} P_h^{\eta} \int_{z*a} \left(\frac{1}{z} + \frac{\zeta^{CBAM}}{a} \right)^{-\eta} \frac{1}{a} d\mu_f(z,a) > 0$$

Consequently, it will increase the domestic prices and therefore reduce domestic demand even further. However, the main aim for the CBAM is not to relieve domestic demand, but to lessen the impact on domestic sales.

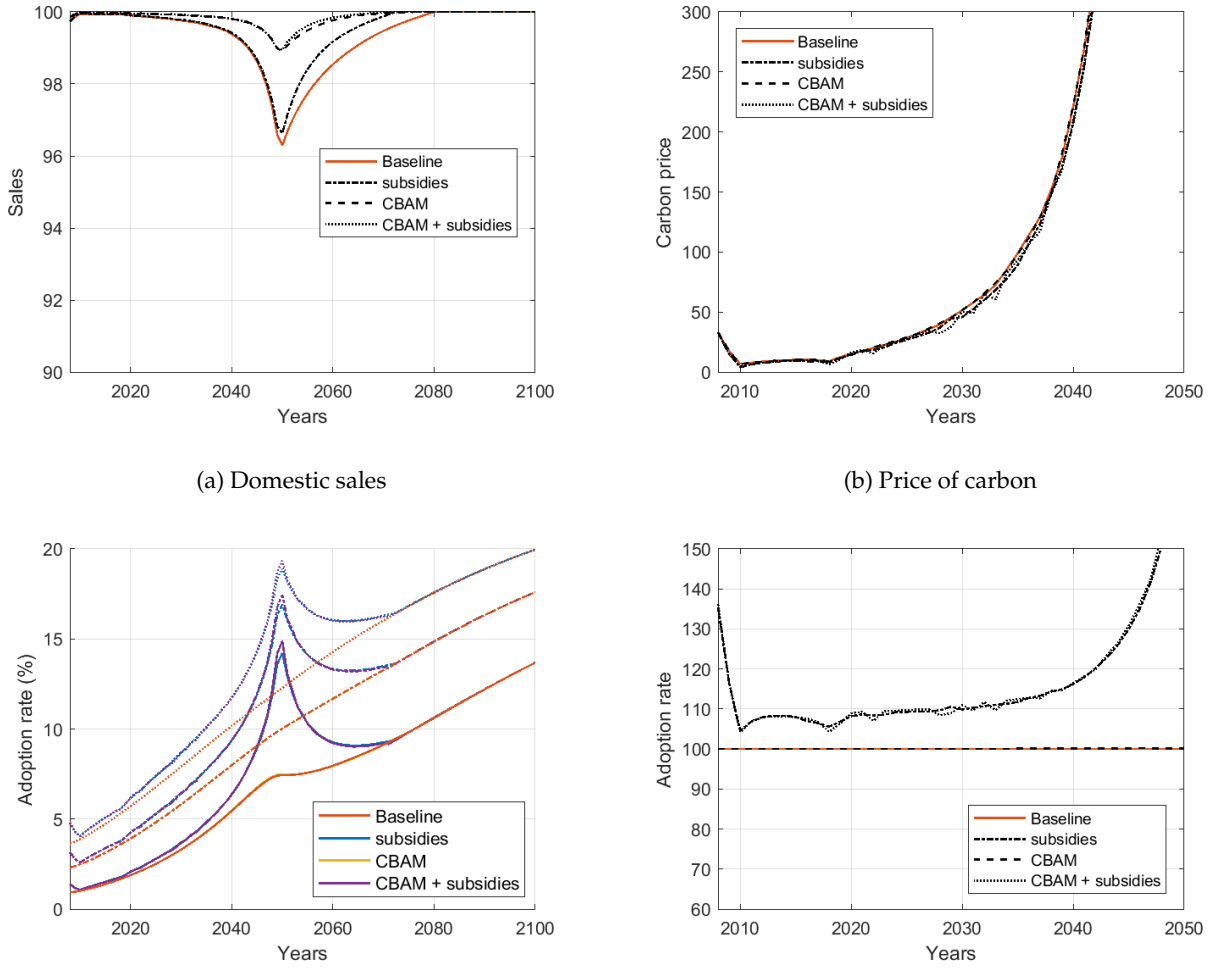
F.2.2 Adoption subsidies

The model now rebates the gains from the taxation lump sum to all firms. Doing this yields zero effect on the firm decisions. One other way to rebate the gains would be to subsidize the adoption. Note that this section still assumes that technological progress behaves at an exogenous rate, which is therefore not affected by the increased demand through the adoption subsidies.

F.2.3 Additional results

Additional results for the various policies are given in Figure F.21. The adoption rate increases significantly when government funds are rebated through subsidies rather than lump sum. The adoption rate is slightly higher for the CBAM than for the baseline scenario, but the implications are limited.

Figure F.21: Additionael Outcome for policy instruments



(a) Domestic sales

(b) Price of carbon

(c) Yearly Adoption rate (%)

(d) Yearly Adoption rate relative to baseline adoption

Note: Evolution of the outcome variables given sensitivity analysis for the 2 policy instruments CBAM and adoption subsidies, as well as their interaction, relative to the baseline model in red.

F.3 Clean technological progress

The baseline model considers clean technological progress as exogenous; it becomes cheaper to adopt novel clean technologies, but at a fixed rate. In this section, I extend this to an endogenous Romer-type model

where the rate of technological progress is a function of the total demand for that technology.

This setting could be further extended to fit into a full directed technological change framework as in Acemoglu et al. (2012), but that falls beyond the scope of this paper.

F.3.1 The innovation sector

Suppose that there is an innovative industry which reacts endogenously on the supply of capital it receives. As long as the market is perfectly competitive - or with CES where markups remain constant - the cost of the good is perfectly proportional to technological progress. Consequently, clean technology maps perfectly in an endogenous Romer-type model. All other equations still hold.

Suppose the productivity of clean technology A_t evolves as a function of the total invested capital $K_t = \sum k_{t,i}^{Q} + k_{i,t}^{F}$ and an idea production function ψ. Not allowing for a scaling parameter as in Acemoglu et al. (2012), the law of motion is $A_t = (1+\psi K_t)A_{t-1}$. In terms of $\gamma_t = 1/A_t$, I rephrase towards:

$$\gamma_t = \frac{1}{A_t} = \frac{1}{(1+\psi K_t)A_{t-1}} = \frac{\gamma_{t-1}}{1+\psi K_t} \tag{F.59}$$

Here, the only exogenously determined variable function is ψ. To obtain ψ, I set up a simple simulated method of moments procedure to match the endogenous gamma for the first 10 periods with the exogenously estimated value $\gamma = 0.9855$ in the baseline model. Formally, I seek the ψ that minimizes the following differential:

$$\min\{abs(\bar{\gamma}^{Exo} - 1/N \sum_{t=1}^{N=11} \gamma_t^{Endo})\} \tag{F.60}$$

The outcome for which this minimization function holds is presented in Table F.15.

Table F.14: Simulated method of moments

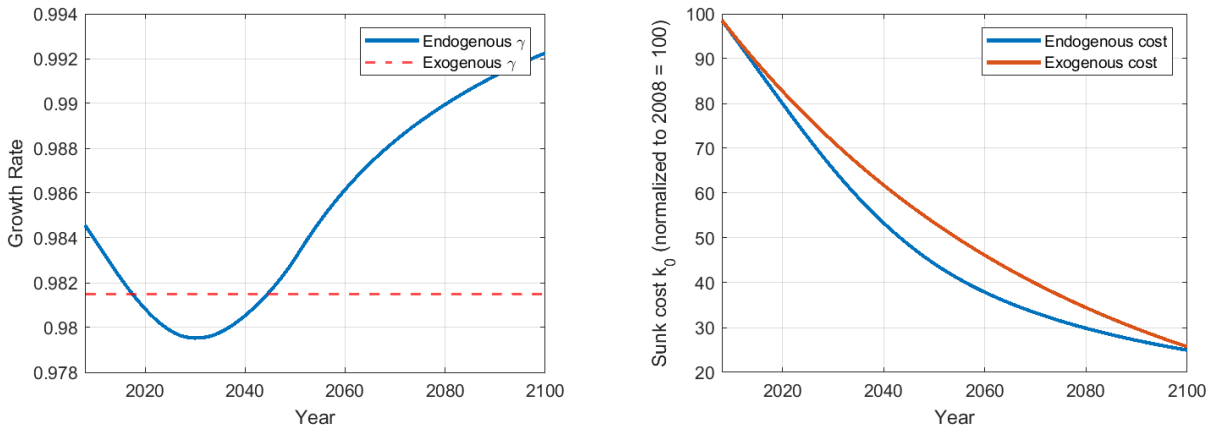
Variable	Value	Moment	Exogenous estimation	Endogenous model
ψ	0.81	$1/N\sum_{t=1}^{N=11}\gamma_t$	0.9855	0.9818

F.3.2 Baseline results

The newly estimated parameter of green technological evolution is depicted in Figure F.22 and compared with the baseline model in the text. Technological progress starts out slower than the exogenous version, but as time progresses, the carbon policy gets more stringent, and more firms feel the need to adopt, the demand rises, increasing the speed at which γ_t depreciates to around 0.978 in 2030. After, both aggregate K diminishes as the price gets cheaper and demand declines. rising the growth rate in γ_t slowly back to higher than the exogenous rate. However, the effect is sufficiently large for the adoption cost k_t to be almost the whole time being structurally lower than the exogenous value (Figure F.22b)

To assess whether the deviation in γ_t has meaningful implications for aggregate output, I examine its effects on three key outcomes: the carbon price, total lost domestic sales relative to the business-as-usual scenario, and the evolution of the adoption rate in both models. Figure F.23 illustrates these results. In the

Figure F.22: Endogenous Technological progress



(a) Growth parameter $\gamma_{i,t}$

(b) Evolution of the cost of abatement

Note: Evolution of the technological progress parameter γ_t in Figure F.22a and the corresponding sunk cost of the adoption decision k_t in Figure F.22b, relative to the baseline exogenous model. The model with the endogenous rate is shown in blue and the exogenous rate is shown in red in both figures. Note that the exogenous rate is constant and does not change over time.

short term (2008 - 2030), aggregate outcomes appear largely similar across the two models. However, as time progresses the pace of technology adoption accelerates, which makes it easier to reach the emissions cap. As a result, a lower carbon price is sufficient to enforce the cap, thereby reducing the impact on aggregate sales.

F.4 Carbon leakage at the extensive margin

In this section, I extend the model to incorporate a reallocation decision. If firms find the policy becoming too stringent, they can opt to leave the domestic economy to a foreign economy where carbon policies are completely absent. While the data rejects this hypothesis : Instances of firm exits are infrequent, and their impact is minimal when they do occur.

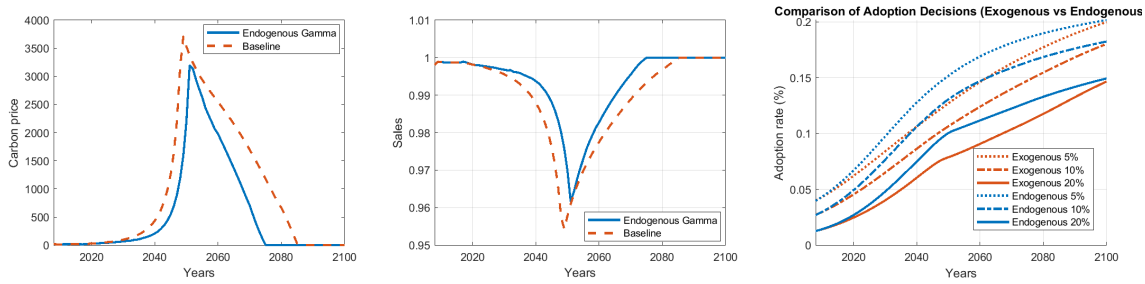
An alternative interpretation is that the extensive margin is effectively muted due to sufficient grandfathering of emission certificates, such that this form of leakage is not observed. To capture this mechanism more explicitly,

F.4.1 The model

I augment the set of actions $\mathscr{D} = \{0, x, \chi_1, \dots \chi_N\}$ with x defined as the relocation decision and the choices in the adoption decision $\chi = \{1.05, 1.10, 1.20\}$. The Bellman equation can be written as:

$$V_t(a_t, z; R_t) = \max_{d_t \in \mathscr{D}} \{\pi(a_{t+1}, z_t, d_t; \Upsilon_{t+1}) + \beta \mathbb{E}[V_{t+1}(a_{t+1}, z_{t+1}; \Upsilon_{t+1} | a_t, z_t, d_t)]\}, \tag{F.61}$$

Figure F.23: Comparison outcome variables with baseline model



(a) Carbon price (b) Sales (c) Adoption rate

Note: Comparison between the baseline model in the text where γ is exogenous (red) and the endogenous model (blue) for three different variables: the price of carbon, aggregate domestic sales and the average adoption rate. Figure F.23a shows the total price of carbon to achieve the carbon cap in either model. Figure F.23b shows the total lost domestic sales relative to the business as usual scenario model between for either model. Figure F.23c depicts the average adoption rate for the three different values $\chi \in \{1.05, 1.10, 120\}$ for the endogenous model (blue) in dotted, dashed and full line respectively and compares them with the exogenous model (red).

with following profit function:

$$\pi(a_{t+1}, z_t, d_t; \Upsilon_{t+1}) = \begin{cases} \pi(a_t, z_t, 0; R_{t+1}) & \text{if } d_t = 0 \\ \pi^f - k_x & \text{if } d_t = x \\ \pi(a_t \chi, z_t, 0; R_{t+1}) - k_{1,t} & \text{if } d_t = \chi \end{cases} \tag{F.62}$$

At the choice $d_t = x$, the firm reaches a terminal state. The value function states $V_t(|d_t = x) \equiv V^f = \pi^f/(1-\beta)$.[30] The foreign profits π^f denote the profit of the firm absent of any carbon policy (i,e. $\pi^f(a_t, z_t) = \pi^f(z_t)$ for all levels of a_t). Note that relative to Equation 1 in the main text, the variable Υ_t replaces ζ_t as this setting takes aggregate fluctuations into account.

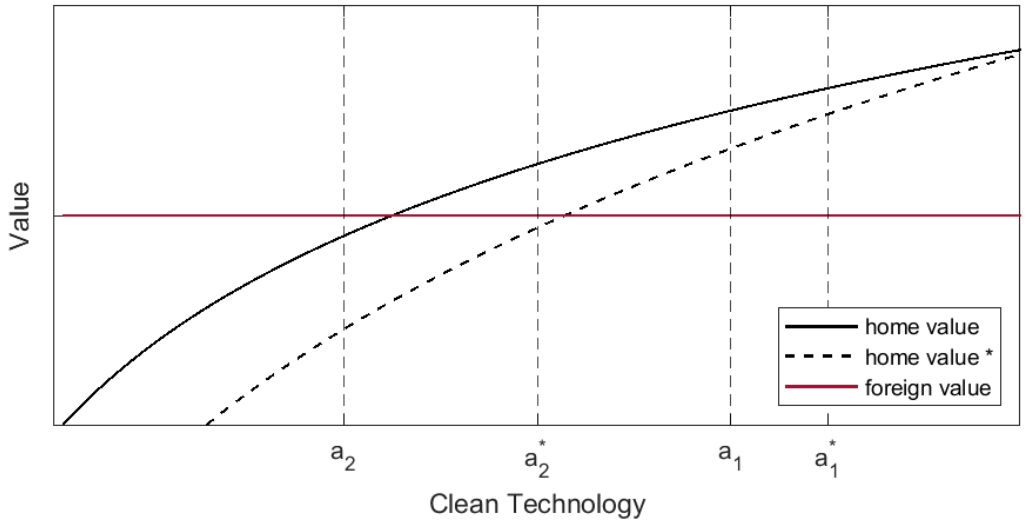
The inclusion of the reallocation generates another layer for policymakers to take into account. Even more than in the baseline setting, a gradual implementation of the carbon policy is required to abide the tradeoff of reduced output relative to reduced emissions. Intuitively, as the policy increases, the decision to reallocate as well.

Adjusting the value of ζ_t would raise the threshold at which firms become indifferent between adopting clean technology but it would also mean that the threshold for firms to exit the domestic market and move towards a pollution heaven increases as well (from $a^{x,L}$ to $a^{x,H}$).

To combat carbon leakage, the government can increase the number of freely allocated certificates κ where $\partial\pi/\partial\kappa < 0$ and $\partial^2\pi/\partial^2\kappa = 0$. Since it does not affect the marginal value of the firm, it does not influence the adoption decision of the firm directly, yet it does affect the decision to move abroad (from $a^{x,2}$

[30] This equality ($V^f = \pi^f/(1-\beta)$) only holds when firms are in partial equilibrium. In the estimated model in Section 4, firms in the foreign country will be affected by the aggregate market price through the relocation decision and the trade margin.

Figure F.24: Increase cost of carbon ζ



to $a^{x,2}$). The optimal carbon policy to ensure that firms do not relocate abroad involves setting the emissions cap such that the value of clean technology of the firm, a, lies between the two threshold values.

Aggregate distribution The number of firms active in the home country is now a function of the decision variables $N_{h,t} = \int \mu_h(z,a_t) \leq 0.5$ and the foreign country $N_{f,t} = \int \mu_f(z,a_t) = 1 - N_{h,t}$. The distribution of firms in the home country exhibits the following motion:

$$\mu_h(z,a_{t+1}) = \int_{z*a} Pr(d_t|a_t,z,\Upsilon_t)F(a_{t+1}|z,a_t,d_t)d\mu_h(z,a_t) \tag{F.63}$$

where the transition operator $F(a_{t+1}|z,a_t,d_t)$ describes the transition law for the firm conditional on remaining at home and $Pr(d_t|a_t,z,\Upsilon_t)$ gives the probability of each decision. If $d_t = 1$, the firm exits the home economy; consequently, $N_{h,t} < N_{h,t_0}$ for all t.

The evolution of the distribution of firms in the foreign country is characterized by an influx of firms from the home country seeking to avoid the carbon policy:

$$\mu_f(z,a_{t+1}) = \mu_f(z,a_t) + \int_{z*a} Pr(d_t = 1|a_t,z,\Upsilon_t)F(a_{t+1}|z,a_t,d_t = 1)d\mu_h(z,a_t)] \tag{F.64}$$

The first part of the equation entails the current distribution of firms currently present in the foreign country. The second part includes the new firms that evade the home country($d_t = 1$) multiplied by the probability of relocating, $Pr(d_t = 1|a_t,z,\Upsilon_t)$, i.e. minus the number of domestic firms. Firms that relocate keep their current level of clean technology a_t and productivity z (Ramondo et al., 2015).

Estimation strategy The model requires the estimation of 1 additional parameter: the cost of relocation for domestic firms k_x. To obtain this value, I match the share of exit among multinational subsidiaries in the data with as a blunt way of accounting for carbon leakage. The number of these subsidiaries is rather small (**??**) however.

Table F.15: Simulated method of moments

Variable	Value	Moment	Data	Model
k_x	0.81	Multinational exit rate	0.007%	0.9818

The indicative results are shown in Table F.16 Note that due to the small sample size - only 12 multinational subsidiaries have opted to close throughout the period from 2013 to 2019 - the error terms are bootstrapped. The results show that among the multinational subsidiaries, the ones that decide to exit the domestic economies are on average more polluting. This holds even after controlling for their level of productivity and their level of freely allocated certificates.

Table F.16: Ex ante characteristics of multinational closures (Phase III)

	(1)	(2)	(3)
VARIABLES	x	x	x
nrl_prf		-0.138**	-0.055
		(0.064)	(0.049)
nrl_ct	-0.103*	-0.084	-0.130*
	(0.054)	(0.057)	(0.079)
ln_ac			-0.008
			(0.006)
Observations	405	396	392
Adjusted R-squared	0.021	0.051	0.034
Country f.e.	✓	✓	✓
Sector f.e.	✓	✓	✓
Bootstrapped s.e.	✓	✓	✓

Note: Relation between the price of carbon certificates and the decision to adopt novel levels of clean technology for more polluting firms. (1) - (3) show the specifications for a 10% increase in clean technology. (1) shows the baseline regression; (2) includes the interaction with an alternate price indicator, the global price of oil; and (3) includes the ratio of fixed certificates to total emissions as a test of the rationality of the firm. Specification (4) and (5) resume the baseline specification but for a reduction of 20% and 30% respectively. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

While the estimates are statistically significant, their interpretation warrants caution. First, significance is only observed at the 10% level. Combined with the small sample size (both in terms of exiting firms and total subsidiaries), it is premature to draw strong conclusions. Ut may be tempting to interpret the results as direct evidence of carbon leakage. However, given the current dataset, it is impossible to deter-

mine whether the capital that would have been allocated to these subsidiaries has instead been redirected toward investments abroad. Further research, using either more granular data or a broader dataset that also includes subsidiaries of multinational firms participating in the ETS outside the European Union, is needed to provide a more robust understanding of the costs and drivers of subsidiary closures among multinationals.

G. Achieving net zero

This section provides further details on the optimization strategy underlying the results presented in Section 6 of the main text. It elaborates on the structure of the optimization problem, defines the policy scenarios considered, and presents the additional steps required to infer the implied social cost of carbon consistent with a net-zero emissions path.

Finally, it presents additional exercises that help to contrast the welfare implications of unilateral versus globally coordinated climate policies.

G.1 The Global 'first-best'

I compare the optimal unilateral scenario with a scenario in which both governments enact the same policy - i.e. a global optimal solution. The latter constitutes a symmetric setting where the distribution in both countries remains exactly the same and balanced trade is preserved. Since income is normalized to $I_h = I_f = 1$, the aggregate sales for either country will sum to 1 as well. Consequently, there is no distortional business-stealing effect between countries. Optimization in this setting aims to search for the b_1 for which the loss in demand is

The government in the unilateral setting has the additional problem that their demand is exposed to a foreign market relative to the 'first-best' solution. Consequently, from the moment they set a policy, the asymmetry between the two countries rises, increasing the share of the demand allocated to the foreign goods; thus reducing the profit margin of the domestic producers. However, as the government is demand-optimizing rather than production-optimizing, the existence of cheaper foreign alternatives is welfare-enhancing.

What is worse relative to the global first-best scenario: as the demand for the foreign goods market rises, their output will increase and therefore foreign emissions. Depending on the assumptions of the trade parameters, even increasing total emissions.[31] The latter is only there as a theoretical scenario, as the iceberg trade costs are there as a mitigating factor to hamper the effect of substitution for aggregate emissions.

[31] As $\tau > 1$ and the economies are symmetric in their emission output such exorbitant negative reallocation will not occur. From the moment $\tau \to 1$ or, what is more likely, the domestic distribution of firms is substantially cleaner than the foreign market, it might happen

G.2 Defining the scenarios

The main scenario discussed in the text adopts a carbon price of $\lambda = 139€/\text{tonne}$, consistent with the social cost of carbon used in the baseline simulations. To evaluate sensitivity, I examine four additional cases that span a wide range of carbon valuations:

- **High valuation:** $\lambda = 1193€/\text{tonne}$. This high valuation case represents a government placing very high value on avoided emissions, thus pursuing aggressive decarbonization. While many consider this still to be extreme, more and more academic research is demonstrating that the social cost of carbon can be valued to higher values that policy makers often consider (Bilal and Kanzig, 2025, 2024).
- **No valuation:** $\lambda = 0€$. This corner case depicts a government that does not internalize the externality of emissions at all. Only from period T onward does the carbon price deviate from zero. In this scenario, the domestic government relies solely on the forward-looking nature of the agents to adjust their clean technology preemptively to achieve net zero by 2050.
- **Moderate valuation:** $\lambda = 42€/\text{tonne}$. This corresponds to a commonly used social cost of carbon estimate in the years prior to 2020 (Acemoglu et al., 2016).
- **Historical valuation:** $\lambda = 4.5€/\text{tonne}$. This value reflects early estimates of the social cost of carbon, as in Nordhaus (1992).

G.3 Optimizing for inferring λ

Finally, drawing upon the insights learned in the previous subsection: there is an optimal λ that minimizes the implied path of carbon.

I first start by searching the value of b_1 that minimizes the differential between $\theta(t)$ and the current cap of carbon using a quadratic loss function for the whole vector from 2008 to 2050. The value for which the differential is minimized is $b_1 = -0.0059$.

The optimization procedure uses finite differences to approximate the first-order condition of the government optimization function in Equation 24. I initiate a trial value of $\lambda = 60€/\text{tonne}$ for which welfare is evaluated at three nearby points: $W(b_1 - \varepsilon, \lambda)$, $W(b_1, \lambda)$, and $W(b_1 + \varepsilon, \lambda)$ with $\varepsilon = 0.0001$. If $W(b_1)$ is greater than both neighbors, then b_1 is locally optimal under that λ. If instead welfare increases when moving right ($W(b_1 + \varepsilon) > W(b)$), the emissions penalty is too low and λ must be increased. If welfare increases when moving left ($W(b_1 - \varepsilon) > W(b)$), the penalty is too high and λ must be reduced. This adjustment is iterated using a binary search.

G.4 The foreign entity

Throughout the analysis, the optimization procedure assumes that the two countries are symmetric, differing only in the presence of a carbon policy. To test the robustness of this assumption, I perform a sensitivity analysis in which the foreign country is assumed to be 10%, 50%, or 100% more polluting than the domestic

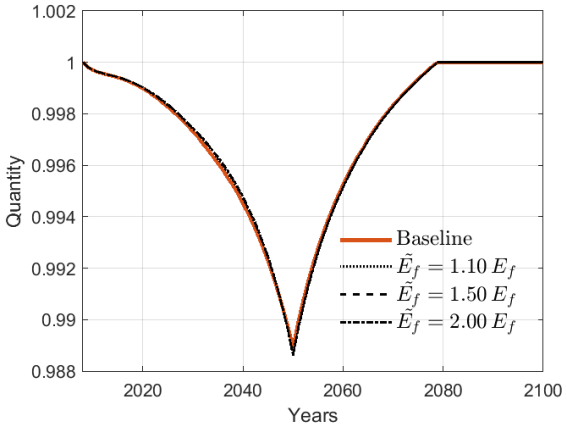
one, while maintaining a carbon price of 139 €/tonne. Table G.17 reports how the optimal curvature parameter b_1 changes as the domestic share in the total amount of emitted emissions declines. As the marginal benefit of the policy decreases, since it affects a smaller portion of total emissions, the government's incentive to reduce domestic emissions weakens, leading to a slower pace of abatement. The results for the

Table G.17: Optimized b_1 for different values of foreign emissions $\tilde{E}_f$

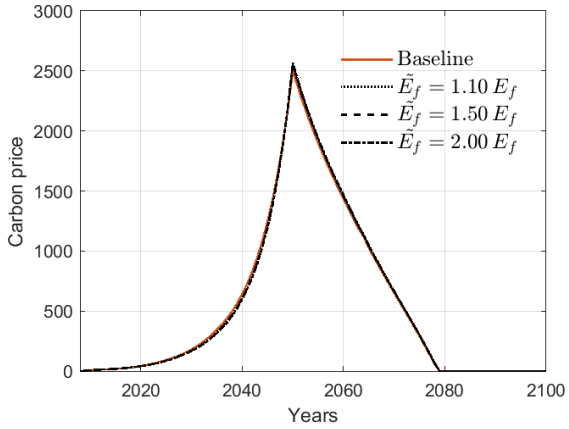
Case	$\tilde{E}_f$	b_1
1.1	$\tilde{E}_f = 1.00E_f$	3.1294×10^{-2}
1.2	$\tilde{E}_f = 1.10E_f$	2.984387×10^{-2}
1.3	$\tilde{E}_f = 1.50E_f$	2.967×10^{-2}
1.4	$\tilde{E}_f = 2.00E_f$	2.927422×10^{-2}

different values of E_f and therefore b_1 are given in Figure G.25. Since the foreign government does not internalize its emissions, there is no feedback from foreign emissions to aggregate prices. The sole departure from the baseline model arises through b_1 in the domestic government's optimization problem and its consequent impact on aggregate prices and firms' idiosyncratic adoption decisions. As shown in the table, the effect remains negligible even for a doubling of emissions.

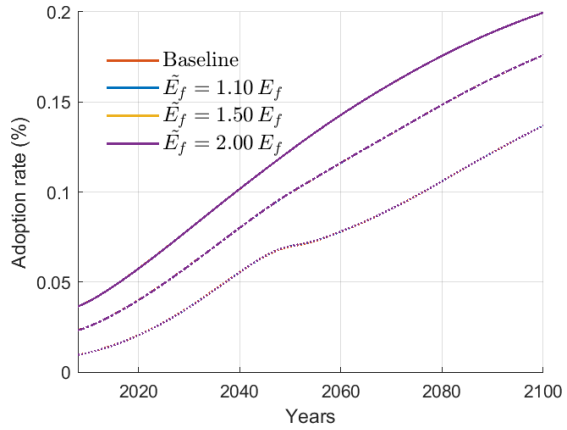
Figure G.25: Implications of increasing importance foreign emissions E_f



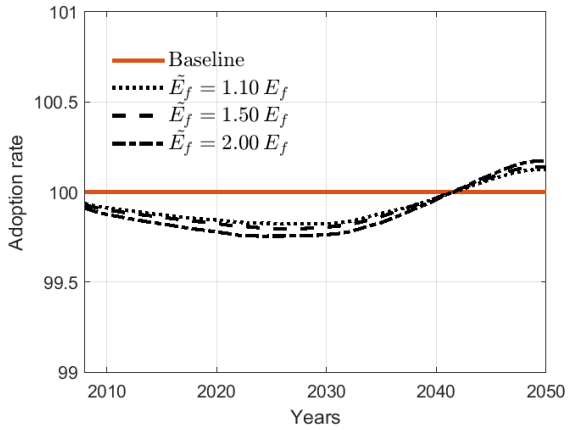
(a) Domestic demand



(b) Price of carbon



(c) Yearly Adoption rate (%)



(d) Yearly Adoption rate (%)

Note: Evolution of the outcome variables given sensitivity analysis for $\tilde{E_f} = \{1, 1.10, 1.50, 2.00\}E_f$. Panel (a) and (b) show the aggregate outcome variables of demand and the price of carbon respectively for the baseline (red) the 5% increase (dotted) and the 5% decrease (dashed). Panel (c) and (d) show the changes in adoption rate for either. the full line represents a $\chi = 1.20$, the dashed shows $\chi = 1.10$ and dotted $\chi = 1.05$.