Pollution Pricing in Equilibrium: Production, Reallocation, and Aggregate Impacts

R. Andrew Butters* Jackson Dorsey† Ivan Rudik‡

February 18, 2026

Abstract

We develop a framework for quantifying the equilibrium effects of pollution pricing using commonly available data on firm financials and emissions. We use the framework to estimate marginal abatement cost curves at the firm, industry, and economy level, which provide a clear decomposition of aggregate emissions reductions into within-firm abatement, reallocation across firms, and technological change. Applying the framework to the European Union Emissions Trading System (2005–2021), we find that firm-level abatement costs are steep in the short run: holding technology fixed, firms cannot easily reduce emissions without cutting output. Yet industry-level and economy-wide abatement cost curves are substantially flatter because production is reallocated toward cleaner firms. We document a 65 percent improvement in clean technology over our sample period, shifting abatement cost curves downward. Together, reallocation and technological change reconcile steep firm-level abatement costs, large aggregate emissions reductions, and small effects on output and employment.

JEL Codes: D22, E23, L11, O33, O44, Q52, Q54, Q58
Keywords: Carbon pricing, marginal abatement costs, EU-ETS, firm heterogeneity, reallocation

*Kelley School of Business, Indiana University. Email: rabutter@iu.edu
†University of Texas at Austin. Email: jackson.dorsey@austin.utexas.edu
‡Cornell University and NBER. Email: irudik@cornell.edu

1 Introduction

Emissions pricing is a central instrument of environmental policy. By attaching a price to pollution, it disincentivizes emissions while preserving the efficiency gains that arise when markets allocate the right to emit. The economic effects of emissions pricing depend critically on how firms and consumers adjust in equilibrium. Firms may respond by altering production processes, investing in cleaner technologies, or reducing output, while consumers shift demand in response to changing prices. Together, these responses determine not only how much emissions fall, but also how the costs of reducing emissions are distributed across firms, industries, and workers. Understanding these equilibrium responses is therefore essential for evaluating the aggregate and distributional consequences of carbon pricing—particularly for identifying which firms and industries are most responsive, and who ultimately bears the costs of pollution abatement.

In this paper, we develop a framework to quantify the equilibrium effects of carbon pricing on emissions and economic outcomes. The model isolates two distinct channels through which carbon pricing reduces aggregate emissions. The first operates within firms: higher pollution prices increase the cost of emissions-intensive production, inducing firms to reallocate inputs toward abatement or adopt cleaner technologies. The second operates across firms: changes in relative output prices shift demand toward less emissions-intensive producers. The strength of these channels—and thus the aggregate cost of emissions reductions—depends on production technologies, demand elasticities, and the distribution of emissions intensity across firms. Our framework provides a unified, empirically tractable approach for estimating these objects and aggregating them to construct marginal abatement cost curves at the firm, industry, and economy-wide levels.

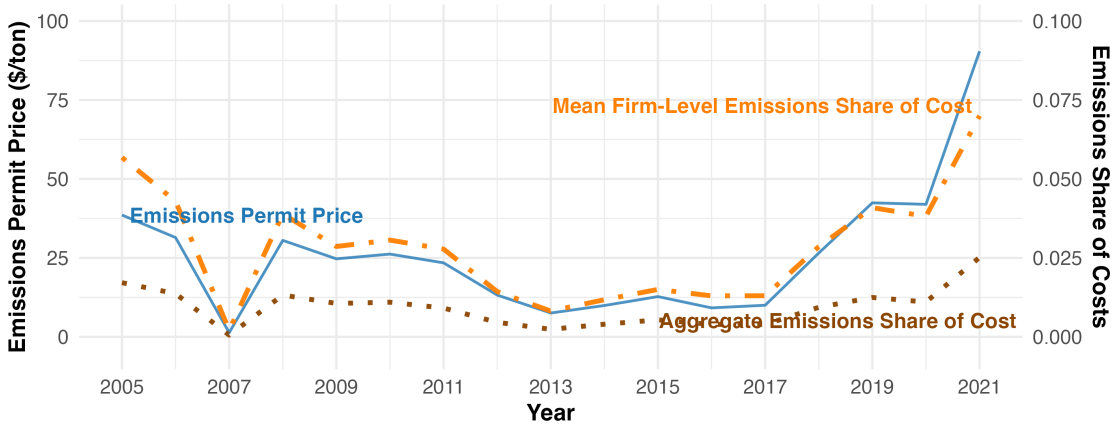
We build the model from microfoundations of firm production and abatement. Firms combine capital, labor, and materials to produce output, generating emissions as a byproduct that can be reduced through costly abatement. A central empirical challenge is that abatement effort is often unobserved: in many settings, researchers observe emissions and input use, but not how resources are allocated to emissions reduction. We show that under constant returns to scale and weak assumptions on the mapping between abatement and emissions, the firm's problem is equivalent to one in which emissions enter directly as an input alongside capital, labor, and materials. This reformulation generalizes the Cobb–Douglas structures in Copeland and Taylor (2004), Shapiro and Walker (2018), and Hollingsworth et al. (2024) to a broad class of production technologies, and it allows estimation of the production function using standard techniques and widely available data on firm financials and emissions. Crucially, the approach does not require auxiliary data on abatement effort

or abatement costs, which are typically unavailable.

We use our model to study the first seventeen years of the European Union Emissions Trading System (EU-ETS), the world's first and largest carbon market. Since its introduction in 2005, EU-ETS has regulated emissions from power generation, manufacturing, and other energy-intensive sectors through a cap-and-trade system. Over our sample period, aggregate emissions at regulated installations declined by more than 30 percent, while industrial production remained relatively stable. This divergence motivates the core policy questions we address: how much of the observed emissions reduction can be attributed to carbon pricing, what are the short- and long-run costs of these reductions, and how do these costs vary across industries?

Our EU-ETS setting highlights the need for a flexible production specification. Figure 1 plots the evolution of EU-ETS permit prices alongside emissions cost shares—defined as a firm's emissions costs (i.e. the economic costs of retired emissions permits) divided by total costs. If firms could easily substitute away from emissions when prices rise, we would expect firm cost shares to rise more slowly than prices as firms reduce emissions to offset the increase. Instead, firm-level emissions cost shares rise nearly in parallel with permit prices, suggesting strong complementarity between emissions and other inputs. Although the firm emissions cost share moves in lockstep with EU-ETS permit prices, the economy-wide aggregate emissions cost share, computed using total emissions and total costs across all firms, tracks it less closely.This divergence between firm-level and aggregate patterns suggests that while individual firms face rigidities in adjusting their emissions intensity, the market as

Figure 1: Permit Prices and Emissions Cost Shares Over Time



Notes: This figure shows the evolution of EU-ETS permit prices alongside emissions cost shares at the firm and aggregate level. The solid line shows the average annual permit price. The dashed line depicts the mean firm-level emissions cost share—defined as a firm's emissions multiplied by the permit price divided by its cost of goods sold. The dotted line shows the aggregate emissions cost share, calculated using total emissions and total cost of goods sold across all firms.

a whole responds through reallocation—shifting production toward less emissions-intensive firms when emissions prices rise.

While our theoretical framework can accommodate a broad class of constant returns to scale production technologies, the patterns in Figure 1 point to two features our empirical specification must capture: near-complementarity between emissions and other inputs at the firm level, and reallocation across heterogeneous firms at the industry level. We show that a nested constant elasticity of substitution (CES) specification fits the data well, and so we adopt this functional form for our baseline analysis. The inner nest combines capital, labor, and materials into a productive composite with a common constant elasticity of substitution between the three. The outer nest combines the productive composite with emissions. The outer nest elasticity governs within-firm substitutability between the composite and emissions and how easily firms can reduce emissions by shifting inputs toward abatement effort. We also allow firms to differ in their *clean technology*—an emissions-augmenting productivity term. Higher clean technology means fewer emissions for any given level of inputs and output, capturing differences across firms in fuel mix, process efficiency, or vintage of capital equipment.[1] An electric utility operating gas-fired plants, for instance, has higher clean technology than one operating coal plants, since gas generation emits roughly half the CO_2 per kilowatt-hour. Similarly, a steel producer using electric arc furnaces has higher clean technology than one using traditional blast furnaces.

A central contribution of our paper is showing how firm-level marginal abatement costs aggregate to industry- and economy-wide marginal abatement cost curves. Even when individual firms face steep abatement costs, aggregate emissions can fall if production shifts from dirtier to cleaner firms. The scope for such reallocation depends on two factors: heterogeneity in emissions intensity and the elasticity of demand. A firm's emissions cost share determines how its unit costs—and therefore its output price—change in response to a carbon price increase; greater heterogeneity means a common carbon price shock generates more dispersion in output prices. The demand elasticity governs how consumers respond to these price changes: when demand is more elastic, market shares shift more aggressively toward cleaner firms. The result is that industry-level marginal abatement cost curves can be substantially flatter than firm-level curves. Our estimation approach closely follows Raval (2019) and Oberfield and Raval (2021), who estimate the capital-labor elasticity of substitution at the firm level and show how firm-level parameters aggregate to industry-level elasticities. We extend their framework to incorporate pollution abatement and show how

[1] In our setting using common firm financial data, we only observe the aggregate level or expenditures on labor, capital, and materials, not whether any given input is emitting, or whether any given input is allocated toward abatement of emissions. In the production function, the clean technology term thus lets us rationalize how the same observed aggregate level of inputs leads to different levels of emissions.

production parameters, demand elasticities, and firm heterogeneity jointly determine the shape of marginal abatement cost curves at each level of aggregation. To our knowledge, we are the first to show how to aggregate firm-level production data to estimate marginal abatement cost curves that account for equilibrium reallocation.

Our estimates reveal four key findings. First, firm-level marginal abatement cost curves are steep across all sectors because the elasticity of substitution between emissions and the productive composite is near-zero. Holding technology fixed, firms have limited ability to reduce emissions by reallocating inputs in the short run—cutting emissions requires cutting output. Within-firm substitution is therefore not a viable margin for substantial emissions reductions.

Second, while steep firm-level abatement costs limit short-run responses given existing technology, we find substantial long-run improvements in clean technology over our sample period.[2] Clean technology reflects changes to production processes that may take years to implement, such as fuel switching or capital upgrades—improvements that shift marginal abatement cost curves downward over time. Aggregate emissions-augmenting productivity increased by approximately 65 percent between 2005 and 2021, meaning that firms require roughly 40 percent fewer emissions to produce the same level of output. This improvement varies considerably across sectors: electricity generation experienced gains exceeding 80 percent, consistent with the shift from coal to natural gas, while sectors such as chemicals showed essentially no technological progress.

Third, industry-level marginal abatement cost curves are substantially flatter than firm-level curves in sectors where demand is elastic and emissions intensity is dispersed. In electricity, petroleum refining, and basic metals—which together account for over 60 percent of emissions in our sample—industries can achieve emissions reductions at far lower cost than individual firms because carbon pricing shifts production toward cleaner producers. Reallocation is the dominant mechanism for emissions reductions in these sectors. The scope for this reallocation is largest precisely where it matters most: in the high-emitting industries that are central to climate policy.

Fourth, counterfactual simulations using our estimated model quantify the equilibrium effects of carbon pricing and decompose the roles of reallocation and technological change. Doubling the permit price from its 2005 level—roughly matching the actual price increase through 2021—reduces aggregate emissions by 37 percent while lowering output by only 2 percent and employment by 1.6 percent. To understand how the economy adjusted so

[2]Formally, clean technology is the change in emissions intensity that remains after accounting for the response to carbon price changes implied by the estimated substitution elasticity. It captures improvements in emissions productivity that are not explained by short-run input reallocation.

smoothly, we simulate the combined effect of the carbon price increase and the observed improvements in clean technology. Together, emissions fall by 46 percent with output declining by less than 1 percent. The additional emissions reductions from technological improvement come with smaller—not larger—output costs, because cleaner technology reduces firms' exposure to carbon prices. These results suggest that reallocation and clean technology have together mitigated the impact of rising carbon prices on economic aggregates like output and employment. Steep firm-level marginal abatement cost curves coexist with flat industry-level curves because reallocation and technological improvement are the primary adjustment margins.

Our framework contributes to several literatures. We build on a line of work that models emissions as an input to production (Copeland and Taylor, 2004; Shapiro and Walker, 2018; Hollingsworth et al., 2024), extending this approach to flexible CES technologies that can accommodate the near-complementarity we find in the data. The methodological approach for estimating substitution elasticities and aggregating them to industry-level parameters follows Raval (2019) and Oberfield and Raval (2021), who developed these techniques for capital-labor substitution; we apply their framework to the emissions margin and show how the resulting aggregate elasticities map to marginal abatement cost curves at different levels of aggregation.

Our work also bridges two strands of the carbon pricing literature that have proceeded largely in parallel. A large applied microeconomic literature uses quasi-experimental approaches to estimate the causal effects of EU-ETS on firm-level outcomes such as emissions, employment, and innovation (Petrick and Wagner, 2014; Calel and Dechezleprêtre, 2016; Marin et al., 2018; Calel, 2020; Dechezleprêtre et al., 2023; Colmer et al., 2024). These studies provide credible and transparent identification but focus on specific subsets of firms and cannot directly speak to aggregate effects or the mechanisms through which firm responses translate to economy-wide outcomes.[3] A separate macroeconomic literature uses aggregate time-series methods to estimate the economy-wide effects of carbon pricing (Metcalf and Stock, 2023; Bjørnland et al., 2024; Känzig, 2025; Ortubai et al., 2025), capturing market-level responses but without decomposing the underlying mechanisms. Our structural approach complements both: we use rich firm-level data to estimate production primitives, then aggregate through our quantified model to recover industry and economy-wide marginal abatement cost curves that reflect both within-firm and reallocation margins.

Finally, although we apply the framework to EU-ETS, the methodology is portable to other settings where firm-level emissions and financial data are available. The framework

[3]Indeed, micro estimates generally difference out important aggregate and equilibrium effects through the use of fixed effects (Nakamura and Steinsson, 2018).

thus provides a template for constructing marginal abatement cost curves and evaluating the aggregate and distributional effects of emissions pricing across a range of regulatory contexts.

The remainder of the paper proceeds as follows. Section 2 describes the EU-ETS institutional setting and our data sources. Section 3 develops the production model and proves the equivalence between the production-abatement formulation and the emissions-as-input representation. Section 4 presents our estimation strategy. Section 5 reports the main parameter estimates and constructs marginal abatement cost curves at the firm, industry, and economy-wide levels. Section 6 presents counterfactual simulations. Section 7 concludes.

2 Institutional Setting and Data

2.1 EU Emissions Trading System

The European Union Emissions Trading System (EU-ETS), established in 2005, is the world's first and largest international carbon market and the cornerstone of EU climate policy. EU-ETS is a cap-and-trade system that sets a total limit on greenhouse gas emissions for regulated sectors and reduces this cap over time. Within this framework, firms receive or purchase emission allowances, each permitting the emission of one metric ton of CO_2. These allowances can be traded, providing economic incentives for companies to reduce their emissions: firms that cut emissions can sell surplus allowances, while those exceeding their limits must buy additional allowances. Compliance is strictly monitored—regulated installations must report their greenhouse gas emissions annually, and these reports are subject to independent third-party verification.

EU-ETS has undergone multiple phases and reforms. Phase 1 (2005–2007) was primarily a "learning phase" to establish baselines and market infrastructure. Phase 2 (2008–2012) coincided with the global financial crisis, which led to lower-than-expected emissions and a large surplus of allowances. This surplus depressed permit prices, weakening incentives for emissions reductions. In response, the EU implemented "backloading"—delaying the auction of 900 million allowances—and subsequently established the Market Stability Reserve (MSR) in 2015, which began operating in 2019. The MSR adjusts allowance supply based on the total number of allowances in circulation, helping to rebalance supply and demand. A 2018 reform strengthened the MSR, contributing to a substantial increase in permit prices that persisted through the end of our sample period. These policy reforms provide a source of exogenous variation in permit prices that we exploit for identification.

In its early years, most allowances were allocated for free, but the system has gradually shifted toward auctioning. The scope of EU-ETS has also expanded over time: while the

program initially covered energy utilities and energy-intensive manufacturing sectors, aviation was added in 2012. We focus our analysis on the set of industries regulated since the program's inception.

2.2 Data Sources

Firm-Level Financial Data. We use data from Bureau van Dijk's Orbis dataset to measure firm-level economic outcomes. Orbis aggregates information from administrative records and private surveys to construct a harmonized dataset for both private and public companies. The dataset includes balance sheet, income statement, and employment data. Variables relevant for our estimation include the cost of labor, fixed assets, employment, and operating revenues. For each firm, Orbis also reports its NACE sector code, country, and address. Orbis has substantial coverage in Europe, and in some countries such as Spain, has close to universal coverage of firms.

Emissions Data. We obtain installation-level emissions from the European Union Transaction Log (EUTL). The EUTL records and tracks permit transactions and verified emissions covered by EU-ETS. Each installation is a physical site regulated by EU-ETS. For each installation-year, the EUTL reports its name and location, verified emissions, and permits allocated, traded, and surrendered.

Crosswalk. We link the Orbis financial data to the EUTL emissions data using the EUTL-Orbis crosswalk provided by the European Commission's Joint Research Centre Data Catalogue (Letout, 2021).

Carbon Prices. We obtain annual average EU-ETS permit prices from the World Carbon Pricing Database maintained by Resources for the Future (RFF). This database compiles carbon pricing data for both cap-and-trade programs and carbon taxes globally. Since the Orbis data are annual, we use the average transaction price for EU-ETS permits as our measure of the emissions price for each year.

2.3 Sample Construction

After merging the datasets and dropping observations with incomplete data, we obtain firm-level financials and verified emissions for over 4,500 firms from 2005–2021 across 22 countries. The matched sample spans a wide range of carbon-intensive sectors such as electricity gener-

ation, cement production, textiles, and steel manufacturing, and covers approximately 50% of verified emissions under EU-ETS.

Appendix Table A1 provides a detailed breakdown of sectoral coverage in our matched sample. Electricity, gas, and steam production dominates the sample, accounting for 43% of total verified emissions and over 1,000 unique firms. Non-metallic mineral products (primarily cement and glass) represents the second-largest emitting sector at 17%, followed by basic metals (10%), refined petroleum products (9%), and chemical products (7%). Together, these five sectors account for approximately 85% of emissions in our sample. The remaining sectors span a diverse range of manufacturing and non-manufacturing activities, from paper and food products to mining and transportation.

Figure A4 in the appendix shows how these sectoral emissions shares have evolved over time. While electricity generation has consistently dominated emissions, its share has declined modestly from 2005 to 2021, reflecting both improvements in power sector efficiency and the relative growth of other regulated sectors. With these data in hand, we now turn to the model that guides our estimation.

3 Model

3.1 Production and Abatement

Firms produce output using capital, labor, and materials, but production generates emissions as a byproduct that can be reduced through costly abatement. The key economic question is how easily firms can substitute away from emissions, which determines how costly abatement is.

The challenge for estimation is that abatement effort is typically unobserved: we see firms' emissions and their use of capital, labor, and materials, but not how they allocate resources toward abatement activities. Our approach addresses this by showing that the underlying production-and-abatement model is *equivalent* to a production function where emissions enter as an input. This equivalence—which we prove holds for a broad class of constant returns to scale technologies in Appendix C—means we can estimate the model using standard production function techniques applied to observed data on outputs, inputs, and emissions, without needing to observe abatement directly.

We now formalize this approach. Production proceeds in two stages: first, firms combine capital, labor, and materials; second, firms choose how much of this productive capacity to allocate toward output versus abatement.

Firm i in sector n at time t produces a productive composite using capital, labor, and

materials:

$$X_{int} = A_{int}\left[\alpha_{K,in}K_{int}^{\frac{\xi_n-1}{\xi_n}} + \alpha_{L,in}L_{int}^{\frac{\xi_n-1}{\xi_n}} + \alpha_{M,in}M_{int}^{\frac{\xi_n-1}{\xi_n}}\right]^{\frac{\xi_n}{\xi_n-1}}. \tag{1}$$

This composite represents potential output before accounting for emissions abatement. A_{int} is Hicks-neutral productivity, the α parameters are share parameters that sum to one across inputs, and ξ_n is the elasticity of substitution between capital K_{int}, labor L_{int}, and materials M_{int}.

Firms abate by diverting a share of their inputs toward abatement activities rather than production. Let $\theta_{int} \in [0, 1)$ denote this abatement intensity—the fraction of the productive composite allocated to abatement rather than output. For example, a steel plant might increase its use of scrap metal relative to virgin ore, which reduces emissions but requires additional labor for sorting and processing, more expenditure on materials to buy different inputs, and greater use of capital equipment for scrap preparation. Conditional on the input and abatement choice, net output is:

$$Y_{int} = (1 - \theta_{int})X_{int}. \tag{2}$$

This formulation is directly analogous to the "share of output devoted to abatement" approach in Copeland and Taylor (2004); Shapiro and Walker (2018); Hollingsworth et al. (2024). Our contribution is to extend this framework from Cobb-Douglas to CES production, and in Appendix C we show how it generalizes to a broader class of constant returns to scale production functions.

The remaining question is how abatement effort translates into emissions reductions. We model this through an abatement function $\psi(\theta_{int})$ that maps abatement intensity to emissions per unit of the productive composite:

$$E_{int} = \psi(\theta_{int}; \pi_{in}, \sigma_n, B_{int})\, X_{int}. \tag{3}$$

The function ψ is decreasing (more abatement means fewer emissions) and convex. We adopt a functional form governed by three parameters:

$$\psi(\theta_{int}) = \frac{1}{B_{int}}\left(\frac{(1-\theta_{int})^{\frac{\sigma_n-1}{\sigma_n}} - (1-\pi_{in})}{\pi_{in}}\right)^{\frac{\sigma_n}{\sigma_n-1}}. \tag{4}$$

Here B_{int} is the firm's *clean technology*—higher B_{int} means fewer emissions for any level of abatement effort. The parameter σ_n governs how quickly marginal abatement costs rise

as the firm abates more. The parameter π_{in} determines the baseline emissions intensity of production, holding technology and substitutability fixed.

The key result is that combining equations (2)–(4) yields an equivalent representation where emissions appear as an input rather than a byproduct:

$$Y_{int} = \left[(1-\pi_{in})X_{int}^{\frac{\sigma_n-1}{\sigma_n}} + \pi_{in}(B_{int}E_{int})^{\frac{\sigma_n-1}{\sigma_n}}\right]^{\frac{\sigma_n}{\sigma_n-1}} \tag{5}$$

In this representation, the same parameters take on a production function interpretation. The parameter σ_n becomes the *elasticity of substitution* between the productive composite and emissions—it determines how easily firms can reduce emissions by reallocating inputs. When σ_n is low, emissions and productive inputs are near-complements, making deep emissions cuts expensive. For example, cement production releases CO_2 directly through calcination—a chemical reaction, not combustion—so no reallocation of inputs can avoid these emissions implying a low σ_n. Electricity generation may allow higher σ_n if plants can shift between fuel sources given existing infrastructure.

The parameter B_{int} is *clean technology*, measured in units of output per emissions. It multiplies emissions in equation (5), so higher B_{int} means a firm can produce more output from each unit of emissions. For example, a firm with gas-fired power plants has higher B_{int} than an otherwise identical firm with coal plants because it emits roughly half the CO_2 per kWh. The share parameter π_{in} determines the baseline emissions intensity of production. Importantly, θ_{int} and B_{int} play distinct roles: θ_{int} is the firm's within-period *choice* of abatement intensity, while B_{int} reflects predetermined differences in installed equipment, process technologies, or fuel mix. Similarly, σ_n is a *short-run* elasticity given current technology; in the longer run, firms may invest to change B_{int}. We are working to extend the model to incorporate endogenous clean technology investment in Section 3.6.

This reformulation is the key to estimation: equation (5) contains only quantities we observe—output, the productive composite, and emissions—even though the underlying model involves unobserved abatement effort θ_{int}. The abatement function $\psi(\cdot)$ has intuitive limiting behavior (emissions reach their unabated maximum when $\theta_{int} \to 0$ and approach zero as $\theta_{int} \to 1$). Figure A3 illustrates how different values of σ affect firms' responses to carbon pricing.

The equivalence between the production-abatement model and the emissions-as-input formulation rests on three assumptions, which we state formally and prove in Appendix C. First, both production and abatement technologies exhibit constant returns to scale, ensuring that abatement intensity depends only on input allocation, not firm size. Second, the abatement function ψ is continuous, strictly decreasing, and strictly convex—emissions

fall with abatement effort, but at an increasing marginal cost. Third, emissions are proportional to potential output and depend on abatement intensity in a scale-invariant way (homogeneous of degree zero in the firm's size). These conditions encompass a broad class of production technologies.

While the equivalence result is general, we adopt the nested CES specification for estimation and counterfactuals. This specification choice aids interpretation but imposes two additional restrictions: (i) constant elasticity of substitution, and (ii) identical substitutability between emissions and each productive input. We test these restrictions empirically in Appendix C by estimating Morishima elasticities—which require no functional form assumptions—between emissions and each input. We find that these elasticities are uniformly low, approximately constant across observed prices, and statistically indistinguishable across input pairs (Table A3 and Figure A6), supporting our use of the nested CES.

3.2 Firm Cost Minimization

Given the production technology, we now characterize firms' cost minimization problem. This delivers the cost function and first-order conditions that determine input demands as functions of prices and technology.

In each period, firm i in sector n is a price-taker in input markets and faces prices P_{int}^K, P_{int}^L, P_{int}^M, and P_t^E for capital, labor, materials, and emissions. Conditional on producing a given quantity Y_{int} and taking input prices and technologies (A_{int}, B_{int}) as given, the firm chooses its inputs to minimize variable costs.[4] Formally, the firm solves:

$$\min_{X_{int},E_{int}} \left\{ P_{int}^X X_{int} + P_t^E E_{int} \;\middle|\; Y_{int} = \left[(1-\pi_{in}) X_{int}^{\frac{\sigma_n - 1}{\sigma_n}} + \pi_{in} (B_{int} E_{int})^{\frac{\sigma_n - 1}{\sigma_n}} \right]^{\frac{\sigma_n}{\sigma_n - 1}} \right\} \quad (6)$$

where P_{int}^X denotes the unit cost of the composite input X_{int}, which is the solution to the analogous inner nest cost minimization problem.

The first-order conditions yield the standard result that the marginal rate of technical substitution equals the price ratio. Given the nested CES functional form, this implies a log-linear relationship between input ratios and relative prices. We express this in terms of

[4] We omit the setup of the firm minimizing the unit cost in the inner nest to produce one unit of X_{int}; however, this follows identically.

the emissions-to-composite ratio, which is the natural measure of emissions intensity:

$$\underbrace{\log \frac{E_{int}}{X_{int}}}_{\text{Emissions Intensity}} = \underbrace{\log\left(\frac{\pi_{in}}{1-\pi_{in}}\right)}_{\text{Constant}} - \sigma_n \underbrace{\log \frac{P_t^E}{P_{int}^X}}_{\substack{\text{Emissions} \\ \text{Composite} \\ \text{Price Ratio}}} + (\sigma_n - 1)\underbrace{\log B_{int}}_{\substack{\text{Clean} \\ \text{Tech}}}. \tag{7}$$

The coefficient $-\sigma_n$ on the emissions price determines how emissions intensity responds to price changes: higher emissions prices induce firms to reduce emissions relative to the productive composite, lowering E/X. When σ_n is low (near-complements), this substitution is limited. We identify σ_n using variation in the emissions price and control for the composite price P_{int}^X in estimation. Cleaner technology (B_{int}) lowers E/X when inputs are complements ($\sigma_n < 1$).

The analogous conditions for the inner nest govern substitution among capital, labor, and materials. With three inputs, any pairwise ratio can be used; we focus on materials relative to labor since they are the more variable inputs:

$$\underbrace{\log \frac{M_{int}}{L_{int}}}_{\text{Input Ratio}} = \underbrace{\log\left(\frac{\omega_{in}}{\delta_{in}}\right)}_{\text{Constant}} - \xi_n \underbrace{\log \frac{P_{int}^M}{W_{int}}}_{\substack{\text{Materials} \\ \text{Wage} \\ \text{Price Ratio}}}. \tag{8}$$

Higher wages induce firms to substitute toward materials, raising M/L; the magnitude of this response identifies ξ_n. We identify ξ_n using variation in wages and control for other input prices P_{int}^K and P_{int}^M in estimation.

3.3 Demand

We complement the supply side with a nested CES demand system under monopolistic competition. This structure allows us to translate changes in marginal costs—induced by carbon pricing or technological progress—into changes in output, prices, and market shares.

A representative consumer has preferences over sectoral aggregates and firm-level varieties:

$$Q_t = \left(\sum_{n \in N} D_{nt}^{\frac{1}{\upsilon}} Q_{nt}^{\frac{\upsilon-1}{\upsilon}}\right)^{\frac{\upsilon}{\upsilon-1}} \qquad Q_{nt} = \left(\sum_{i \in I_n} D_{int}^{\frac{1}{\epsilon_n}} Q_{int}^{\frac{\epsilon_n-1}{\epsilon_n}}\right)^{\frac{\epsilon_n}{\epsilon_n-1}} \tag{9}$$

where Q_{int} is the quantity demanded of firm i in industry n at time t, D_{int} is a taste or quality shifter, and ϵ_n and υ denote the elasticities of substitution across varieties within a sector and across sectors.

Given these preferences, firm i faces an isoelastic demand curve:

$$Q_{int} = Q_{nt} \left(\frac{P_{int}}{P_{nt}} \right)^{-\epsilon_n}, \tag{10}$$

where P_{int} is the price charged by firm i in industry n and $P_{nt} = \left(\sum_{i \in I_n} D_{int} P_{int}^{1-\epsilon_n} \right)^{\frac{1}{1-\epsilon_n}}$ is the sectoral price index.

Facing CES demand, monopolistically competitive firms optimally set prices as a constant markup μ_n over marginal cost:

$$\mu_n = \frac{P_{int}}{MC_{int}} = \frac{\epsilon_n}{\epsilon_n - 1}. \tag{11}$$

The firm's problem can thus be decomposed into two stages: first, minimize costs conditional on output (Section 3.2); second, set price as a markup over marginal cost and let demand determine quantity.[5]

3.4 Equilibrium

An equilibrium in each period t consists of firm-level prices, quantities, and input choices

$$\{P_{int}, Y_{int}, X_{int}, E_{int}, K_{int}, L_{int}, M_{int}\}_{i,n}$$

together with sectoral aggregates $\{Q_{nt}, P_{nt}\}_n$, such that each firm i in sector n minimizes costs (Section 3.2), sets prices as a markup over marginal cost (Section 3.3), and markets clear:

$$Y_{int} = Q_{int} \qquad Q_{nt} = \sum_{i \in I_n} Q_{int}.$$

We assume input markets are perfectly competitive: capital, labor, materials, and emissions are supplied perfectly elastically at exogenous prices $(P_{int}^K, P_{int}^L, P_{int}^M, P_t^E)$, which firms take as given. In estimation, input prices are observed in the data. For counterfactual simulations, this assumption implies that factor prices are unaffected by policy-induced changes in quantities, allowing us to isolate the mechanisms operating through product markets and within-firm substitution. The emissions price P_t^E is the policy variable we vary in counterfactuals.

[5]Because the markup is a scalar wedge between the output price and marginal cost, it does not affect the conditional cost-minimizing input mix.

3.5 Marginal Abatement Costs

Our model allows us to derive closed-form marginal abatement cost schedules at three levels of aggregation: the firm, the industry, and the regulated market as a whole. The curvature of these schedules—how rapidly costs rise as emissions fall—depends on within firm abatement potential and reallocation opportunities across firms. We develop each level in turn.

3.5.1 Firm-Level MAC

The marginal abatement cost measures how much it costs a firm to reduce emissions by one unit while holding output fixed. Formally, let $C_{int}(Y_{int}, E_{int})$ denote the minimum cost of producing output Y_{int} with emissions level E_{int}. The firm's MAC is:

$$MAC_{int} \equiv -\frac{\partial C_{int}}{\partial E_{int}}. \tag{12}$$

Given our nested CES production structure, Appendix Section F shows that the firm-level MAC takes the form:

$$MAC_{int}(Y_{int}, E_{int}) = \underbrace{MC_{int}(Y_{int}, E_{int})}_{\substack{\text{marginal cost} \\ \text{of output}}} \cdot \underbrace{A_{int}^{\frac{\sigma_n - 1}{\sigma_n}} \pi_{in} Y_{int}^{\frac{1}{\sigma_n}}}_{\substack{\text{output scale} \\ \text{and technology}}} \cdot \underbrace{B_{int}^{\frac{\sigma_n - 1}{\sigma_n}}}_{\substack{\text{clean} \\ \text{technology}}} \cdot \underbrace{E_{int}^{-\frac{1}{\sigma_n}}}_{\substack{\text{slope of} \\ \text{MAC curve}}} \tag{13}$$

with an elasticity with respect to emissions of $-\frac{1}{\sigma_n(1-s_{E,int})}$, reflecting the fact that reducing emissions at fixed output raises marginal cost as the firm substitutes toward the productive composite.[6]

The elasticity $-\frac{1}{\sigma_n(1-s_{E,int})}$ implies the MAC curve slopes upward in abatement: the cost of abatement rises as emissions fall. The steepness of this curve is governed by $1/\sigma_n$. When emissions and productive inputs are strong complements and σ_n is small, the MAC curve is steep, reflecting the difficulty of cutting emissions without cutting output. This connects to the evidence in Figure 1: firm-level emissions cost shares track permit prices closely, suggesting limited substitutability and therefore steep firm-level MACs. However, σ_n governs how the emissions-composite *mix* responds to an emissions price change through cutting emissions and increasing the composite, not the emissions level. The $1 - s_{E,int}$ term adjusts this ratio effect so that it yields the effect on emissions. For more emissions intensive firms, the change in the mix is driven primarily by ramping up the composite instead of cutting back on emissions. The more emissions intensive a firm is, the more the emissions-composite elasticity must be shaded down to obtain the level emissions response to a change

[6] Although output is held fixed, marginal cost increases as emissions fall because the firm must re-optimize its input mix away from emissions and toward other productive inputs.

in the emissions price.

Clean technology B_{int} enters with exponent $(\sigma_n - 1)/\sigma_n$, which is negative when $\sigma_n < 1$. Improvements in clean technology therefore shift the MAC curve downward, reducing the cost of achieving any given level of abatement. A gas-fired power plant, for instance, has higher B_{int} than an otherwise identical coal plant, implying a lower MAC at any emissions level.

At the firm's cost-minimizing input choices, the MAC admits a particularly intuitive representation (Appendix F):

$$MAC_{int}(Y_{int}, E_{int}) = MC^*_{int}(Y_{int}, E_{int}) \cdot s^*_{E,int}(Y_{int}, E_{int}) \cdot \frac{Y_{int}}{E_{int}} \tag{14}$$

where $s^*_{E,int} \equiv \frac{P^E_t E_{int}}{P^E_t E_{int} + P^X_{int} X^*_{int}}$ denotes the emissions cost share evaluated at the firm's cost-minimizing input choices. In equilibrium, all firms facing the common permit price P^E_t choose emissions such that their MAC equals the permit price: $MAC_{int} = P^E_t$ for all i. Equation (14) therefore characterizes an equilibrium relationship—firms with higher unit costs must have lower emissions cost shares or lower emissions intensity to satisfy this condition. A cleaner firm (higher B_{int}) has both lower $s_{E,int}$ and lower Y_{int}/E_{int}, which is how it reaches the common equilibrium MAC despite having a lower MAC curve.

3.5.2 Industry-Level MAC

We define the industry MAC analogously: the minimum cost of reducing aggregate industry emissions by one unit while holding industry output fixed. Let $C_{nt}(Y_{nt}, E_{nt})$ denote the minimum cost of producing industry output Y_{nt} subject to aggregate emissions E_{nt}. The industry MAC is

$$MAC_{nt} \equiv -\frac{\partial C_{nt}}{\partial E_{nt}}. \tag{15}$$

This cost function reflects optimal allocation of production and emissions across firms within the industry—not just the sum of firm-level costs.

An industry can therefore reduce emissions more cheaply than the sum of its firms would suggest. Beyond within-firm substitution along individual MAC curves, the industry has an additional margin: reallocation of output from high-emission to low-emission producers. The curvature of the industry MAC depends on two forces that determine how effectively this reallocation channel operates.

The first is heterogeneity in emissions intensity. When firms differ in their emissions cost shares, carbon pricing changes their relative costs—dirtier firms become relatively more expensive. The wider the distribution of emissions intensities within an industry, the larger

the potential gains from shifting output toward cleaner producers. Figure A1 documents substantial heterogeneity in emissions cost shares across firms, particularly in sectors like petroleum products and non-metallic minerals, suggesting significant scope for reallocation-based abatement.

We summarize this heterogeneity with an index χ_n:

$$\chi_n = \frac{\sum_i (s_{E,int} - \Gamma_n)^2 s_{int}}{\Gamma_n(1-\Gamma_n)} \tag{16}$$

where Γ_n is the industry's aggregate emissions cost share and s_{int} is firm i's share of industry costs. The index is bounded between zero and one. It equals zero when all firms have identical emissions intensities, leaving no scope for reallocation. It approaches one when emissions intensity is maximally dispersed—some firms are perfectly clean while others are highly polluting.

The second force is demand elasticity. Reallocation requires consumers to shift purchases from dirty to clean firms as relative prices change. The within-sector demand elasticity ε_n governs how readily this occurs. When ε_n is high, products are close substitutes and consumers easily redirect spending toward cleaner firms whose prices have risen less. When ε_n is low, products are differentiated and consumers incur substantial welfare costs from changing consumption patterns, limiting the extent of reallocation.

Following Oberfield and Raval (2021), we show in Appendix G that these two forces combine into an aggregate elasticity of substitution that governs the industry-level response to carbon pricing:

$$\bar{\sigma}_n = \underbrace{(1-\chi_n)\sigma_n}_{\text{within-firm}} + \underbrace{\chi_n \varepsilon_n}_{\text{reallocation}} \tag{17}$$

The aggregate elasticity is a weighted average of the micro elasticity σ_n, which governs within-firm substitution, and the demand elasticity ε_n, which governs between-firm reallocation. The weight on the reallocation channel is the heterogeneity index χ_n.

The industry MAC curve cannot be solved in closed form, however the slope is the industry aggregate analogue of the firm slope and is given by: $-\frac{1}{\bar{\sigma}_n(1-\Gamma_n)}$. This expression makes transparent how the three key parameters—σ_n, ε_n, and χ_n—shape the industry MAC. The slope $-\frac{1}{\bar{\sigma}_n(1-\Gamma_n)}$ is decreasing in all three since $\bar{\sigma}_n$ is increasing in all three if $\varepsilon_n > \sigma_n$ as we will see in our results. Higher within-firm substitutability (σ_n) flattens the curve by making it easier for each firm to abate. Higher demand elasticity (ε_n) flattens the curve by making it easier for consumers to switch to cleaner firms. Higher heterogeneity (χ_n) flattens the curve by increasing the weight on the typically larger demand elasticity and by expanding the scope for cost-saving reallocation.

When firms differ in emissions intensity ($\chi_n > 0$) and demand is elastic ($\varepsilon_n > \sigma_n$), the aggregate elasticity exceeds the micro elasticity: $\bar{\sigma}_n > \sigma_n$. The industry MAC is therefore flatter than firm-level MACs for firms near or above the industry average emissions cost share, however the industry MAC may be steeper than firm MACs for firms that are the cleanest in the industry. The industry will typically achieve emissions reductions at lower cost because it can reallocate output toward cleaner firms rather than forcing each firm up its own steep MAC curve. When firms are homogeneous ($\chi_n \approx 0$), only within-firm substitution operates and $\bar{\sigma}_n \approx \sigma_n$.

3.5.3 Aggregate MAC

The aggregate MAC is defined analogously at the market level. Let $C_t(Y_t, E_t)$ denote the minimum cost of producing aggregate output Y_t subject to total emissions E_t, where the cost function reflects optimal allocation across sectors. The aggregate MAC is

$$MAC_t \equiv -\frac{\partial C_t}{\partial E_t}. \tag{18}$$

The same aggregation logic that operated within industries now extends across sectors. At the market level, emissions reductions occur through three nested channels: within-firm input substitution, within-industry reallocation across firms, and across-industry reallocation of final demand. The last channel operates through the across-sector demand elasticity υ: when carbon pricing raises costs more in emissions-intensive sectors, consumers shift expenditure toward goods from cleaner sectors. The strength of this channel depends on how substitutable goods are across sectors and how much sectors differ in their aggregate emissions intensity.

Paralleling the within-industry aggregation, we define a cross-sector heterogeneity index χ_t^{agg} that measures dispersion in emissions cost shares across sectors, weighted by expenditure (Appendix H). The market-level elasticity takes the familiar form:

$$\bar{\sigma} = (1 - \chi_t^{\text{agg}})\bar{\sigma}^N + \chi_t^{\text{agg}}\upsilon \tag{19}$$

where $\bar{\sigma}_t^N$ is a weighted average of sectoral elasticities. The aggregate MAC curve slope is given by: $-\frac{1}{\bar{\sigma}(1-\Gamma)}$ where Γ is the economy-wide emissions cost share.

The comparative statics mirror those at the industry level. Higher within-sector elasticities ($\bar{\sigma}^N$) flatten the aggregate MAC by reducing abatement costs within each sector. Higher cross-sector demand elasticity (υ) flattens the curve by enabling consumers to shift toward cleaner sectors. Higher cross-sector heterogeneity (χ_t^{agg}) flattens the curve by expanding the

scope for cross-sector reallocation and by placing more weight on the demand elasticity υ. More emissions intensive the economy, the higher Γ and the more difficult it is to reduce the level of emissions relative to changing the emissions-composite mix.

This three-level structure—firm to industry to market—reveals how reallocation progressively flattens the MAC schedule at each stage of aggregation. Even when firm-level MACs are steep because σ_n is small, elastic demand combined with heterogeneous emissions intensities can generate substantially flatter industry and market-level MACs. The cost of achieving aggregate emissions reductions through carbon pricing may therefore be considerably lower than firm-level MACs would suggest.

3.6 Extension: Endogenous Clean Technology Investment

The static model treats clean technology B_{int} as predetermined, which is appropriate for short-run counterfactuals where capital stocks are fixed. In ongoing work, we are developing a dynamic extension that endogenizes B through investment in clean capital. The key innovation is to decompose the firm's capital stock into dirty capital K_d (high output per dollar, high emissions) and clean capital K_c (lower output per dollar, low emissions), with clean technology determined by the clean capital share $\phi \equiv K_c/K$ and the technology frontier Ω_t: specifically, $B_{int} = \Omega_t \cdot \phi_{int}^{\gamma}$. Firms then face a dynamic tradeoff: investing in clean capital today reduces effective output capacity (since clean capital is less productive, $\lambda < 1$) but raises future B, lowering abatement costs and emissions tax liability in subsequent periods. This framework enables several analyses beyond our current setup. First, it allows us to decompose observed improvements in B into exogenous frontier shifts (Ω_t rising as clean technologies improve industry-wide) versus endogenous portfolio shifts (ϕ rising as firms respond to carbon pricing by investing in cleaner capital). Second, it provides a foundation for long-run counterfactuals—such as evaluating 2030 or 2050 climate targets—where capital stocks adjust to policy. Third, it enables comparison of carbon pricing with alternative policies such as clean investment subsidies. Finally, by inverting our recovered $\hat{B}_{int}$ panel using the structural relationship, we can back out the unobserved clean capital shares and estimate how investment composition responds to carbon price changes, providing microfounded evidence on the dynamic margins of adjustment to environmental policy.

Using the recovered B_{int} panel, we can estimate the elasticity of clean technology with respect to the emissions price: $\frac{\partial \log B_{int}}{\partial \log P_t^E}$. Since clean technology investment is a longer-run process than input substitution, we use local projections to trace out the impulse response of B_{int} to carbon price shocks, instrumenting with the Känzig (2025) surprises.

4 Estimation

Taking the model to data requires estimating production elasticities, ξ_n and σ_n, demand elasticities, ϵ_n and υ, as well as recovering firm-level clean technology B_{int} and constructing the heterogeneity index χ_n that governs reallocation. We proceed sequentially. First, we estimate the inner-nest elasticity ξ_n between capital, labor, and materials using a granular instrumental variable strategy for wages; this yields the composite input price P^X_{int} needed for subsequent steps. Second, we estimate the outer-nest elasticity σ_n between the productive composite and emissions, instrumenting for permit prices using regulatory surprises. The residuals from this regression capture firm-specific variation in clean technology, which we invert to recover B_{int}. Third, we show that the heterogeneity index χ_n—which determines the scope for reallocation—can be computed directly from the cross-sectional dispersion in emissions cost shares, providing a transparent link between the outer-nest residuals and the aggregate elasticity of substitution. Finally, we estimate within-sector demand elasticities ϵ_n from observed markups and the cross-sector elasticity υ from industry-level price-quantity variation in the EU KLEMS database. Together, these estimates allow us to fully quantify our model as well as construct firm-, industry-, and economy-wide marginal abatement cost curves.

4.1 Production Function: Capital-Labor-Materials Elasticity

We estimate the inner-nest elasticity ξ_n using conditional factor demands. Under CES production, the optimal ratio of any two inputs j and k satisfies an analogous relationship with the common elasticity parameter ξ_n—a restriction we test empirically in Section 4.3. For concreteness, consider the materials-to-labor ratio:

$$\log \frac{M_{irnt}}{L_{irnt}} = \xi_n \log P^L_{rt} + \beta_K \log P^K_{irnt} + \beta_M \log P^M_{irnt} + h_{1,n}(t) + \lambda_{1,irn} + \varepsilon_{1,irnt}, \qquad (20)$$

where $h_{1,n}(t)$ is a sector-specific flexible time trend, $\lambda_{1,irn}$ is a firm fixed effect, and $\varepsilon_{1,irnt}$ is the error term. We estimate this regression separately for each sector n, and introduce a region subscript r because our instrument construction depends on a region-year notion of labor markets. The coefficient ξ_n on the log wage identifies the elasticity of substitution: when wages rise, firms substitute toward materials, and the magnitude of this substitution reveals ξ_n.[7]

The CES structure implies that we can also estimate this elasticity using the capital-

[7]Under CES, the coefficient on $\log P^M$ could also identify ξ_n. We focus on $\log P^L$ because our granular instrumental variable strategy provides a credible instrument for wages but not for materials prices.

to-labor ratio, $\log(K/L)$, which yields the same ξ_n under the null hypothesis. In practice, we estimate both specifications and report results for both the M/L and K/L ratios, which provides a specification check on the CES nesting structure.

The empirical challenge is that wages are endogenous: local wages may comove with unobserved productivity shocks that also affect firms' input choices. We address this using a Granular Instrumental Variable (GIV) following Gabaix and Koijen (2024). The key insight is that idiosyncratic labor demand shocks to large firms in a local labor market act as residual labor supply shifters for other firms. When a dominant employer expands, it bids up local wages for all firms in the region, providing exogenous wage variation that is uncorrelated with any individual firm's productivity.

Specifically, we define labor markets at the region-year level and construct a leave-one-out instrument that aggregates idiosyncratic employment shocks to other firms, weighting by their employment shares. Because the firm size distribution in European labor markets is concentrated, shocks to large employers can generate substantial market-wide wage movements. We extract regional wages from firm-level data using a fixed effects decomposition that separates market-level wage variation from firm-specific factors. Technical details of the GIV construction are provided in Appendix E.

The inner-nest elasticity estimates yield sector-specific ξ_n, which we use to construct the composite input price index P^X_{int} and quantity X_{int} for subsequent estimation of the outer nest.

4.2 Production Function: Emissions Elasticity

We estimate the outer-nest elasticity σ_n between emissions and the productive composite using an analogous conditional factor demand approach. From the first-order conditions, the optimal emissions-to-composite ratio satisfies:

$$\log \frac{E_{int}}{X_{int}} = -\sigma_n \log P^E_t + \beta_X \log P^X_{int} + \underbrace{h_{2,n}(t) + \lambda_{2,in} + \varepsilon_{2,int}}_{\propto \log B_{int} \text{ (Clean Technology)}}, \tag{21}$$

where $h_{2,n}(t)$ is a sector-specific flexible time trend, $\lambda_{2,in}$ is a firm fixed effect, and $\varepsilon_{2,int}$ is the error term. We estimate this regression separately for each sector n. The coefficient $-\sigma_n$ on the log emissions price identifies the elasticity: when permit prices rise, firms reduce emissions relative to the composite, and the magnitude of this response reveals σ_n.[8]

A key identification challenge is that the emissions price P^E_t has no cross-sectional

[8]Under CES, the coefficient on $\log P^X$ could also identify σ_n. We focus on $\log P^E$ because our identification strategy instruments for the emissions permit price using high-frequency regulatory surprises.

variation—all firms face the same permit price in a given year—so we cannot include time fixed effects. This raises the concern that unobserved changes in clean technology may be correlated with permit prices. When firms improve their clean technology, their production processes become less emissions-intensive, reducing aggregate permit demand and driving down equilibrium permit prices. This generates a negative correlation between $\log B_{int}$ and $\log P_t^E$ that would bias OLS estimates of σ_n toward zero. The sector-specific flexible time trends $h_{2,n}(t)$ absorb smooth technological progress, but concerns remain if technology and prices comove at higher frequencies.

Our identification strategy isolates exogenous variation in permit prices driven by discrete regulatory events. We instrument for the emissions price using the high-frequency carbon price surprise series developed by Känzig (2025) and extended by Jezequel-Royer and Levieuge (2025). The intuition is as follows: Känzig (2025) constructs a narrative event list of EU-ETS regulatory announcements that affect the expected supply of allowances—European Commission decisions on allowance auctioning, European Parliament votes on market stability reserves, and court judgments on emissions caps. He then measures the change in carbon futures prices in a tight window around each announcement, isolating the "surprise" component attributable to regulatory news rather than slow-moving macroeconomic conditions. These high-frequency price movements around discrete policy events are plausibly uncorrelated with firms' technological choices, which evolve over longer horizons.

Jezequel-Royer and Levieuge (2025) extend this surprise series through 2023 using the same methodology. We aggregate the daily surprises to annual frequency by summing within each calendar year. The resulting annual instrument captures the cumulative regulatory news affecting permit prices in each year, purged of confounding macroeconomic, financial, and energy market trends. The flexible time trends absorb common technological progress, while the instrument isolates year-to-year price variation driven by regulatory surprises rather than correlated technology shocks.

Given the inner-nest elasticity estimates ξ_n, we construct the composite input price P_{int}^X as a CES price index over capital, labor, and materials (see Appendix D). We observe expenditures on capital, labor, and materials separately in the data, which we aggregate to obtain total expenditure on productive inputs $P_{int}^X X_{int}$ and thereby construct X_{int}. We control for $\log P_{int}^X$ but do not interpret its coefficient β_X causally.

The regression structure in equation (21) provides a direct link to clean technology. The firm fixed effects $\lambda_{2,in}$ absorb time-invariant differences in B_{in} (e.g., a gas plant versus a coal plant), while the sector trends $h_{2,n}(t)$ capture industry-wide technological progress. The residual $\varepsilon_{2,int}$ therefore reflects firm-specific, time-varying deviations in clean technology not explained by prices or common trends. As we show in Section 4.5, this residual variation

determines the scope for carbon pricing to reduce emissions through reallocation across heterogeneous firms.

4.3 Validating the Nested CES Specification

The nested CES specification assumes that emissions are equally substitutable with all productive inputs. That is, when the emissions price rises, firms should reduce emissions relative to capital, labor, and materials by the same proportion—the substitution response should not depend on which input we compare emissions to. This restriction justifies treating capital, labor, and materials as a composite in the outer nest.

We test this assumption by estimating Morishima elasticities of substitution, which measure how each input ratio j/E responds to a change in the emissions price, for $j \in \{K, L, M\}$. Under our nested CES, all three elasticities should be identical and equal to σ_n. We estimate each elasticity separately using regressions of $\log(j/E)$ on $\log P^E$, controlling for other input prices and including firm fixed effects and flexible time trends. We instrument for the carbon price using the same regulatory surprises as in Section 4.2. Appendix C.4 provides full estimation details.

Table A3 reports the results. All three elasticities are uniformly low—below 0.1—and statistically indistinguishable from each other, supporting the nested CES restriction. We also test whether the elasticity is constant across observed prices by including a quadratic term $(\log P^E)^2$; the coefficients are small and insignificant (Table A4), consistent with CES.

4.4 Clean Technology Recovery

The regression structure in equation (21) provides a natural decomposition of clean technology B_{int}. Comparing the theoretical model (equation 7) to our estimating equation, we see that clean technology enters through the regression components not explained by prices:

$$(\sigma_n - 1) \log B_{int} = \lambda_{2,in} + h_{2,n}(t) + \varepsilon_{2,int}. \tag{22}$$

The firm fixed effect $\lambda_{2,in}$ captures time-invariant technology differences (e.g., a gas plant versus a coal plant), the sector trend $h_{2,n}(t)$ captures industry-wide technological progress, and the residual $\varepsilon_{2,int}$ captures firm-specific, time-varying technology shocks. Rearranging:

$$\log B_{int} = \frac{1}{\sigma_n - 1} \left(\lambda_{2,in} + h_{2,n}(t) + \varepsilon_{2,int} \right). \tag{23}$$

Alternatively, we can recover B_{int} directly by inverting the first-order conditions. Ex-

pressing objects in terms of the emissions cost share $s_{int}^E \equiv P_t^E E_{int}/(P_t^E E_{int} + P_{int}^X X_{int})$:

$$B_{int} = \left(\frac{1 - s_{int}^E}{s_{int}^E}\right)^{\frac{1}{1-\sigma_n}} \left(\frac{\pi_{in}}{1-\pi_{in}}\right)^{\frac{\sigma_n}{1-\sigma_n}} \frac{P_t^E}{P_{int}^X}. \tag{24}$$

Both approaches yield the same B_{int} panel given the estimated σ_n.

4.5 Emissions Heterogeneity Index

The heterogeneity index χ_n is critical in determining the scope for reallocation when carbon prices change. If firms within an industry differ substantially in their emissions intensity, a carbon price increase raises costs more for dirty firms than clean ones, shifting market shares toward cleaner producers and reducing industry emissions even if no individual firm changes its technology. The heterogeneity index χ_n captures the extent of this cross-firm dispersion and determines the weight on reallocation in the aggregate elasticity (equation 17). We now show that χ_n can be computed from the cross-sectional distribution of emissions cost shares, which connects directly to the outer-nest residuals.

We compute χ_n directly from observed emissions cost shares. Let $s_{E,int} \equiv P_t^E E_{int}/(P_t^E E_{int} + P_{int}^X X_{int})$ denote firm i's emissions cost share—the fraction of firm i's total costs spent on emissions permits. Let $s_{int} \equiv (P_t^E E_{int} + P_{int}^X X_{int})/\sum_j (P_t^E E_{jnt} + P_{jnt}^X X_{jnt})$ denote firm i's share of total industry costs. The industry-level emissions cost share is then:

$$\Gamma_n \equiv \sum_i s_{E,int} \cdot s_{int},$$

which is a cost-weighted average of firm-level emissions cost shares. Intuitively, Γ_n measures the fraction of industry-wide costs accounted for by emissions permits; it is high in emissions-intensive industries and low in clean industries. We define the heterogeneity index as:

$$\chi_n = \frac{\sum_i (s_{E,int} - \Gamma_n)^2 \, s_{int}}{\Gamma_n(1-\Gamma_n)}. \tag{25}$$

The numerator is the cost-weighted variance of firm-level emissions cost shares around the industry mean Γ_n. The denominator $\Gamma_n(1-\Gamma_n)$ normalizes this variance by its maximum possible value (which occurs when firms are split between $s_{E,int} = 0$ and $s_{E,int} = 1$). Thus $\chi_n \in [0, 1]$: when all firms have identical emissions intensities, $\chi_n = 0$; when intensity is maximally dispersed, χ_n approaches one.

We now establish the explicit link between χ_n and the objects estimated in Section 4.2.

The outer-nest regression decomposes clean technology into estimated components:

$$\log B_{int} = \frac{1}{\sigma_n - 1}\left(\hat{\lambda}_{2,in} + \hat{h}_{2,n}(t) + \hat{\varepsilon}_{2,int}\right), \tag{26}$$

where $\hat{\lambda}_{2,in}$ is the estimated firm fixed effect, $\hat{h}_{2,n}(t)$ is the estimated sector time trend, and $\hat{\varepsilon}_{2,int}$ is the residual. At any point in time t, cross-sectional variation in $\log B_{int}$ across firms within sector n comes from two sources: the firm fixed effects $\hat{\lambda}_{2,in}$ (persistent technology differences) and the residuals $\hat{\varepsilon}_{2,int}$ (transitory technology shocks). Denoting this cross-sectional variance as $\mathrm{Var}_i^s(\cdot) \equiv \sum_i (\cdot - \bar{\cdot})^2 s_{int}$:

$$\mathrm{Var}_i^s\left(\hat{\lambda}_{2,in} + \hat{\varepsilon}_{2,int}\right) = (\sigma_n - 1)^2 \,\mathrm{Var}_i^s(\log B_{int}). \tag{27}$$

The connection to χ_n comes through emissions cost shares. From the first-order condition in equation (7), the emissions cost share satisfies:

$$\frac{1}{s_{E,int}} = 1 + \left(\frac{1-\pi_{in}}{\pi_{in}}\right)^{\sigma_n} \left(\frac{P_{int}^X}{P_t^E}\right)^{1-\sigma_n} B_{int}^{\sigma_n - 1}. \tag{28}$$

At a given t, all firms face the same P_t^E, so cross-sectional variation in $s_{E,int}$ is driven by variation in B_{int} (and P_{int}^X, π_{in}). Applying a first-order approximation around the industry mean:

$$s_{E,int} - \Gamma_n \approx \Gamma_n(1-\Gamma_n)(\sigma_n - 1)(\log B_{int} - \overline{\log B}_n). \tag{29}$$

Substituting into equation (25) and using equation (27):

$$\chi_n \approx (\sigma_n - 1)^2 \,\Gamma_n(1-\Gamma_n)\,\mathrm{Var}_i^s(\log B_{int}) = \Gamma_n(1-\Gamma_n)\,\mathrm{Var}_i^s\left(\hat{\lambda}_{2,in} + \hat{\varepsilon}_{2,int}\right). \tag{30}$$

Equation (30) is the key result: the heterogeneity index χ_n can be computed directly from the cross-sectional dispersion of the outer-nest regression's firm effects and residuals. Industries where the estimated $\hat{\lambda}_{2,in}$ and $\hat{\varepsilon}_{2,int}$ are tightly clustered have low χ_n and limited scope for reallocation. Industries where they are widely dispersed have high χ_n, and carbon pricing can reduce emissions by shifting output toward cleaner producers.

This connection clarifies why χ_n enters the aggregate elasticity formula $\bar{\sigma}_n = (1-\chi_n)\sigma_n + \chi_n\epsilon_n$. When χ_n is high, firms differ substantially in their emissions intensity, so carbon pricing creates large differences in relative costs. With elastic demand (ϵ_n large), consumers shift toward cleaner firms, flattening the industry MAC. When χ_n is low, firms are homogeneous and the industry MAC is governed by within-firm substitution (σ_n).

4.6 Demand Elasticity

We estimate industry-level demand elasticities ϵ_n using observed markups. Equation (11) shows that under CES demand and monopolistic competition, the markup is inversely related to the demand elasticity.

We compute markups μ_{nt} at the sector-year level by dividing total revenue by total costs across firms in each sector-year:

$$\mu_{nt} = \frac{\sum_{i \in n} R_{int}}{\sum_{i \in n} c_{int}},$$

where R_{it} is firm i's revenue and c_{it} is total expenditures on labor, capital, materials, and emissions permits. Given the CES structure, we then back out the demand elasticity as $\epsilon_{nt} = \mu_{nt}/(\mu_{nt} - 1)$.

We estimate the cross-industry demand elasticity υ using data from EU KLEMS, a sectoral database that provides industry-level price and quantity indices for European countries constructed from national accounts.[9] We regress the KLEMS sectoral quantity indices on the KLEMS sectoral price indices:

$$\log Q_{nt} = -\upsilon \log P_{nt} + h_n(t) + \varepsilon_{3,nt}. \tag{31}$$

To identify υ we instrument for $\log P_{nt}$ using sector n's average product of labor. Higher labor productivity shifts industry supply; for widely traded products or small industries, this should not meaningfully shift demand through income effects.

With the demand elasticities in hand, we can compute the aggregate elasticity $\bar{\sigma}_n = (1 - \chi_n)\sigma_n + \chi_n \epsilon_n$ for each sector, which governs the slope of industry-level MAC curves. The cross-sector elasticity υ plays an analogous role at the economy-wide level: it determines how output reallocates across industries when carbon pricing raises costs differentially across sectors. Just as within-industry reallocation flattens sector MACs when ϵ_n is large, cross-industry reallocation flattens the aggregate MAC when υ is large. Together, ϵ_n and υ govern the scope for demand-side responses to carbon pricing at the industry and economy-wide levels, respectively.

5 Results

This section presents the estimates of the key model parameters. We organize the findings to mirror the estimation discussion in Section 4: we first present the elasticity of substitu-

[9] EU KLEMS (Capital, Labor, Energy, Materials, Services) provides harmonized productivity and growth accounts across EU member states. See https://euklems-intanprod-llee.luiss.it/ for documentation.

tion estimates for both the inner and outer nests of the production function, then describe the recovered clean technology trends, followed by the heterogeneity indices and demand elasticities that determine the scope for reallocation. We conclude by constructing marginal abatement cost curves at the firm, industry, and economy-wide levels. The final section validates our nested CES specification and investigates robustness of the parameter estimates.

5.1 Elasticities of Substitution

Our nested CES production structure requires estimating two elasticities: the inner-nest elasticity ξ_n governing substitution among capital, labor, and materials, and the outer-nest elasticity σ_n governing substitution between the productive composite and emissions. We discuss each in turn.

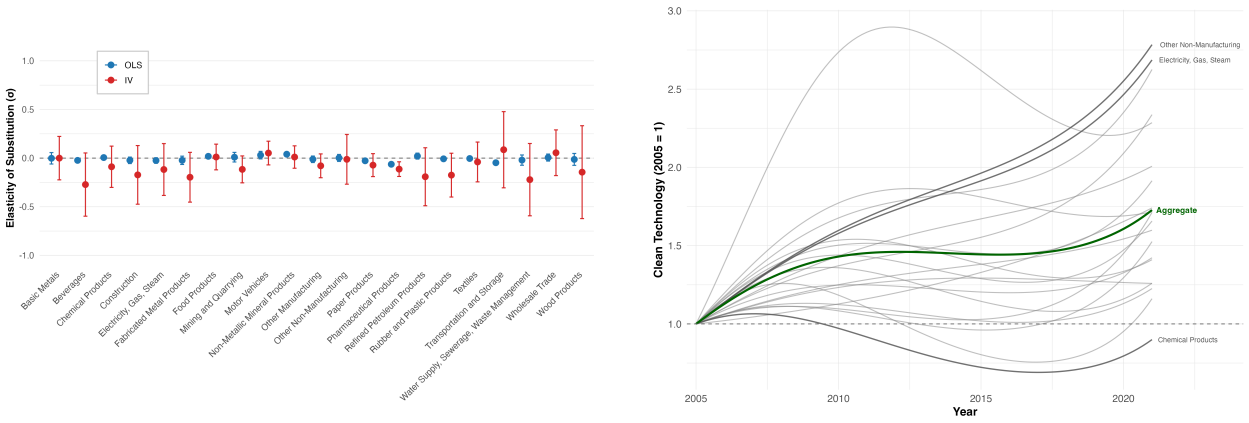
Inner Nest: Capital-Labor-Materials. Figure A2 in Appendix A presents sector-specific estimates of ξ_n from equation (20). Unlike the emissions elasticity discussed below, the inner-nest elasticity varies substantially across sectors. OLS estimates cluster around 1, suggesting approximately Cobb-Douglas substitution among productive inputs. IV estimates show more dispersion: some sectors (electricity, mining) exhibit elasticities above 1, indicating that capital, labor, and materials are gross substitutes, while others (basic metals, fabricated metals) show estimates around 0.5, indicating less flexibility in input mix.

These inner-nest elasticities affect our analysis primarily through their role in constructing the composite input price P^X_{int} and quantity X_{int}, which enter the outer-nest regression. The variation in ξ_n across sectors implies different aggregation weights when combining capital, labor, and materials into the productive composite.

Outer Nest: Emissions-Composite. The central parameter governing firm-level marginal abatement costs is the elasticity of substitution σ_n between emissions and the productive composite. Panel (a) of Figure 2 presents our sector-specific estimates from equation (21), with OLS estimates in blue and IV estimates (instrumenting with carbon price surprises) in red.

The results are striking: all estimates cluster near zero, regardless of sector or estimation method. The OLS and IV point estimates are statistically indistinguishable and economically similar across all industries. We can reject Cobb-Douglas ($\sigma = 1$) for every sector, but cannot reject Leontief ($\sigma = 0$) for any sector. This near-zero elasticity implies that firms have extremely limited ability to substitute away from emissions toward other inputs in the short run—emissions and productive inputs are near-complements in production.

Figure 2: Emissions Elasticity and Clean Technology Trends



(a) Emissions-Composite Elasticity (σ_n)

(b) Clean Technology Trends

Notes: Panel (a) plots sector-specific estimates of the elasticity of substitution σ_n between emissions and the productive composite from equation (21). Blue circles show OLS estimates; red diamonds show IV estimates using carbon price surprises from Känzig (2025) and Jezequel-Royer and Levieuge (2025). Error bars represent 95% confidence intervals. The dashed vertical line marks zero. Panel (b) plots the estimated sector-specific clean technology trends recovered from the cubic time series in equation (21). Each curve represents the level of clean technology B_{nt} for a sector, normalized to 1 in 2005. The thick green line shows the emissions-weighted aggregate trend.

The contrast between the two nests is instructive. Firms have moderate flexibility in combining capital, labor, and materials (ξ_n near or above 1), but almost no flexibility in trading off the productive composite against emissions ($\sigma_n \approx 0$). This pattern reflects the technological rigidity of emissions: while firms can adjust their input mix, the emissions associated with a given level of output are largely determined by installed capital and fuel type rather than by short-run operational choices. As we show in Section 5.4, the near-zero emissions elasticity implies steep firm-level MAC curves, meaning that within-firm abatement—absent technological change—is prohibitively costly.

5.2 Clean Technology Trends

Given the central role of clean technology in shifting MAC curves, we now present the recovered technology trends from our estimation. Panel (b) of Figure 2 plots the sector-specific cubic time trends $h_{2,n}(t)$ from equation (21), which capture industry-wide improvements in clean technology over our sample period.

The aggregate trend, shown by the thick green line, reveals substantial improvement in clean technology over 2005–2021. Starting from a normalized value of 1 in 2005, aggregate clean technology increased to approximately 1.65 by 2021—a 65% improvement. Intuitively,

this means that emissions have become 65% more productive: holding other inputs fixed, firms require approximately $1 - 1/1.65 \approx 40\%$ fewer emissions to produce the same level of output. This aggregate improvement masks pronounced sectoral heterogeneity. The electricity sector experienced one of the largest gains, with clean technology more than doubling over the sample period. This dramatic improvement reflects the well-documented shift from coal to natural gas and renewable generation in European electricity markets. At the other extreme, the chemicals sector shows essentially no technological progress, with the trend remaining flat throughout the sample period.

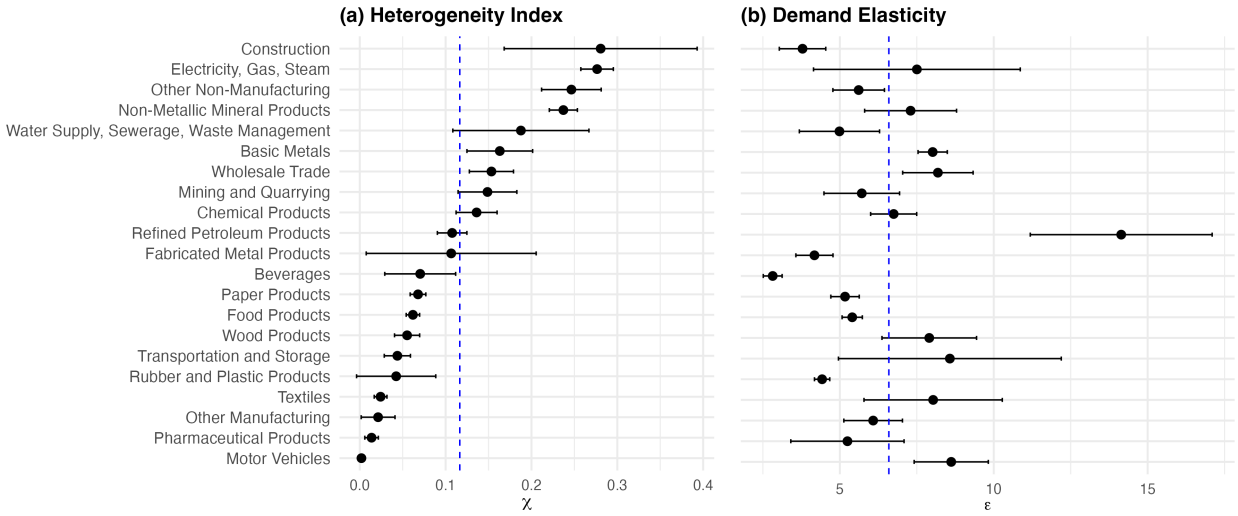
It is important to distinguish the low substitution elasticities from the clean technology trends. The near-zero elasticity σ_n captures firms' short-run abatement possibilities given their existing production technology—how easily they can reduce emissions by reallocating inputs while holding their capital stock and production processes fixed. Clean technology, by contrast, reflects long-run changes to the production technology itself, such as the adoption of new equipment or processes that may take years to build or implement. Changes in emissions intensity over time that are not explained by responses to short-run carbon price shocks—for example, fuel switching or capital upgrades—appear in our framework as improvements in clean technology. While we cannot attribute these improvements to carbon pricing versus other factors such as fossil fuel prices, technological spillovers, or regulatory mandates, the magnitude and heterogeneity of these trends establish that technology adoption—not short-run input substitution—has been the primary margin through which firms have reduced emissions intensity over our sample period.

5.3 Heterogeneity Index and Demand Elasticities

The slope of industry-level abatement curves depends not only on firm-level substitution elasticities but also on the scope for reallocation across heterogeneous firms. Two parameters govern this reallocation channel: the heterogeneity index χ_n, which captures the dispersion in emissions intensity across firms within a sector, and the demand elasticity ϵ_n, which determines how readily consumers shift toward cleaner producers when carbon pricing raises dirty firms' costs. Figure 3 presents both estimates.

Panel (a) reveals substantial variation in emissions heterogeneity across sectors. The electricity, gas, and steam sector exhibits the highest heterogeneity index at approximately 0.27, reflecting the wide dispersion in emissions intensity across firms. Appendix Figure A1 illustrates this heterogeneity, showing the distribution of emissions cost shares across firms within the highest-emitting industries. By contrast, sectors like paper and chemicals have heterogeneity indices below 0.10, indicating that firms within these industries are more ho-

Figure 3: Heterogeneity Index and Demand Elasticities



Notes: Panel (a) shows the heterogeneity index χ_n measuring dispersion in firm-level emissions cost shares within each sector, computed following equation (25). Standard errors are obtained via bootstrap resampling. Panel (b) shows industry demand elasticities ϵ_n estimated from observed markups: $\epsilon_n = \mu_n/(\mu_n - 1)$, where μ_n is the sector-level markup. Standard errors are computed using the delta method. Error bars represent 95% confidence intervals. Sectors are ordered by heterogeneity index.

mogeneous in their emissions profiles. The magnitude of these heterogeneity indices matters for the aggregate elasticity formula $\bar{\sigma}_n = (1 - \chi_n)\sigma_n + \chi_n\epsilon_n$: with $\sigma_n \approx 0$, the first term vanishes and the industry-level elasticity is approximately $\bar{\sigma}_n \approx \chi_n\epsilon_n$, meaning that nearly all of the industry's emissions response to carbon pricing operates through reallocation rather than within-firm substitution.

Panel (b) reports the estimated demand elasticities, which vary by a factor of six across sectors. Petroleum products exhibits the highest demand elasticity at over 14, reflecting the commodity nature of refined fuels and consumers' ability to substitute across suppliers. Basic metals and electricity also show high demand elasticities above 5. At the other extreme, fabricated metals and non-metallic minerals have demand elasticities below 3, consistent with more differentiated products and less consumer flexibility. These demand elasticities interact with heterogeneity to determine the overall scope for reallocation: sectors with both high χ_n and high ϵ_n can achieve substantial emissions reductions through shifting output to cleaner firms, even when firm-level substitution possibilities are negligible. For example, consider the electricity sector with $\chi_n \approx 0.27$, $\epsilon_n \approx 7.5$, and $\sigma_n \approx 0$. The aggregate elasticity is $\bar{\sigma}_n \approx 0.27 \times 7.5 \approx 2$.

Cross-Industry Demand Elasticity. We also estimate the elasticity of substitution across sectors, υ, which governs reallocation of final demand from emissions-intensive to

cleaner industries. Table A2 in Appendix A reports the results. The OLS estimates in columns (1)–(2) are approximately -0.4, suggesting inelastic demand across sectors. However, these estimates are likely biased toward zero due to the simultaneity of supply and demand: positive productivity shocks raise quantities and lower prices, generating a spurious positive correlation that attenuates the estimated price elasticity.

To address this, we instrument for sectoral prices using average labor productivity, which shifts supply but—for traded goods and small industries—should not directly affect demand. The IV estimates in columns (3)–(4) are substantially larger in magnitude, ranging from -2.8 to -3.2, indicating elastic cross-sector demand. The preferred specification in column (4), which includes NACE section-by-year trends to absorb broad sectoral shocks, yields an estimate of $\upsilon \approx 3.2$. This implies that consumers are willing to substitute across industries when relative prices change, providing an additional margin for emissions reductions beyond within-industry reallocation. When carbon pricing raises costs more in emissions-intensive sectors, consumers shift expenditure toward cleaner alternatives, flattening the economy-wide marginal abatement cost curve relative to the industry-level curves.

5.4 Marginal Abatement Cost Curves

The elasticity estimates in the previous subsection imply that marginal abatement cost (MAC) curves at the firm level are nearly vertical when output is held fixed. With an elasticity of substitution between emissions and the productive composite close to zero, individual firms have very limited scope to reduce emissions through within-firm input substitution while holding output fixed. As a result, achieving even modest reductions in emissions at the firm level would require large increases in marginal abatement costs.

We now synthesize the estimated parameters to construct marginal abatement cost curves at higher levels of aggregation, incorporating reallocation across producers and sectors. Figure 4 presents industry-level and economy-wide MAC curves on a log–log scale, where differences in slope directly reflect differences in the effective elasticities governing emissions responses once reallocation is incorporated.

Several patterns emerge. First, industry-level MAC curves are clearly not vertical, although they remain steep. This indicates that emissions reductions become feasible at the industry level once reallocation across producers is incorporated, but that marginal abatement costs still rise quickly as emissions fall. This flattening quantifies the reallocation mechanism: once reallocation is incorporated, industries can achieve emissions reductions more cheaply than any individual firm by shifting output across heterogeneous producers. The degree of flattening varies across sectors and reflects differences in the product $\chi_n \epsilon_n$,

Figure 4: Industry MAC Curves and The Role of Clean Technology



Notes: The left panel plots all 2005 industry aggregate MACs. The right panel plots the economy-wide aggregate MAC with the 2005 and 2021 levels of clean technology. All curves are displayed in log–log space. The industry and economy-wide aggregates are constructed by simulating our model on a grid of counterfactual emissions prices and recovering the corresponding industry and economy-wide aggregate emissions.

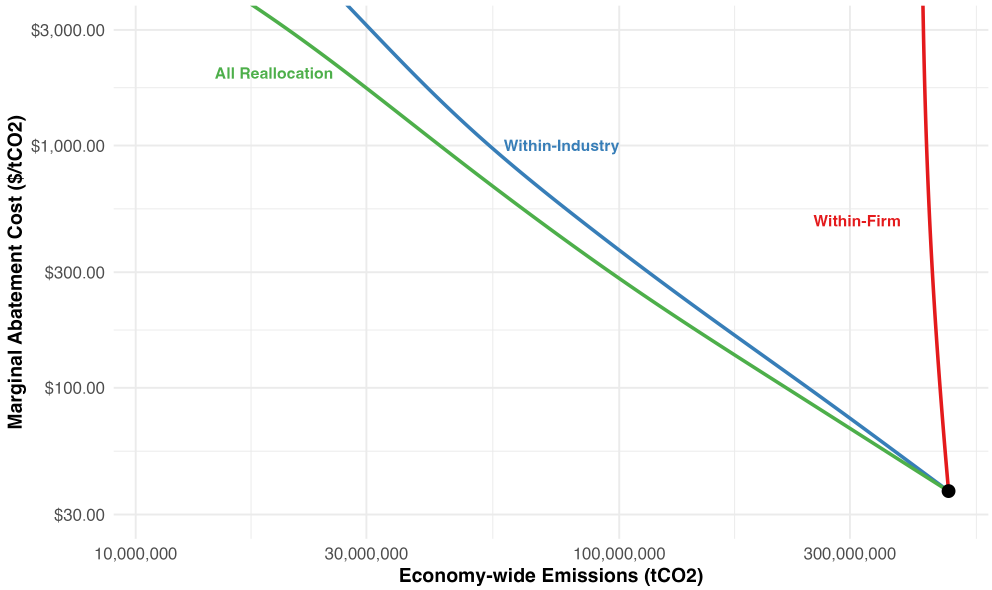
which governs the strength of within-industry reallocation through emissions heterogeneity and demand elasticity.

Non-metallic mineral products are among the industries with visibly flatter MAC curves. This pattern reflects a combination of substantial heterogeneity in emissions intensity across producers and an above-average elasticity of demand, which together allow output to reallocate toward relatively cleaner producers as carbon prices rise. The interaction of emissions heterogeneity and elastic demand generates meaningful within-industry reallocation, producing a noticeably flatter industry-level MAC despite the overall steepness of the curve.

The right panel of Figure 4 plots the economy-wide MAC under the 2005 and 2021 levels of clean technology. Improvements in clean technology shift the aggregate MAC outward, lowering the marginal cost of achieving any given level of emissions. The slope of the economy-wide MAC is similar across the two years, indicating that clean technology primarily shifts the level of abatement costs rather than altering their curvature.

Figure 5 isolates the contribution of each reallocation margin to the aggregate MAC. The red curve shows the economy-wide MAC when only within-firm substitution is permitted—that is, when output shares across firms and industries are held fixed. This curve is extremely steep, reflecting the near-zero firm-level elasticities. The blue curve adds within-industry reallocation, allowing output to shift across firms within each sector while holding cross-sector shares fixed. This curve is substantially flatter, demonstrating that reallocation across heterogeneous firms within industries is quantitatively important for reducing abatement costs. Finally, the green curve allows for full reallocation both within and across industries, matching the economy-wide MAC in Figure 4. The additional flattening from across-industry

Figure 5: Economy-Wide MAC Curves Under Alternative Reallocation Mechanisms



Notes: This figure plots the economy-wide 2005 MAC curve under different reallocation mechanisms. The red line incorporates only within-firm adjustment. The blue line incorporates reallocation across firms within industries. The green line incorporates full reallocation and matches the economy-wide MAC curve in Figure 4.

reallocation is modest but non-negligible, indicating that cross-sector substitution provides a meaningful additional margin for cost-effective abatement.

These patterns directly reflect our estimated parameters and the aggregation formulas derived in Section 3.5. With $\sigma_n \approx 0$, the aggregate elasticity formula $\bar{\sigma}_n = (1-\chi_n)\sigma_n + \chi_n\epsilon_n$ from equation (17) reduces to approximately $\bar{\sigma}_n \approx \chi_n\epsilon_n$. For electricity, with $\chi_n \approx 0.27$ and $\epsilon_n \approx 7.5$, this yields an industry-level elasticity of approximately 2—substantially larger than the near-zero firm-level elasticity and implying a much flatter industry MAC curve. Sectors with lower heterogeneity or less elastic demand, such as chemicals ($\chi_n \approx 0.08$), have smaller aggregate elasticities and steeper industry MACs closer to firm-level curves. At the economy-wide level, the cross-sector demand elasticity $\upsilon \approx 3.2$ provides an additional reallocation margin through equation (19). The right panel of Figure 4 also illustrates the role of clean technology: the 65% improvement in B between 2005 and 2021 shifts the economy-wide MAC curve substantially to the right, reducing the cost of achieving any given level of emissions.

5.5 Robustness: Validating the Nested CES Specification

Our analysis relies on a nested CES structure that treats capital, labor, and materials as a composite in the outer nest with emissions. This specification assumes that emissions are equally substitutable with all productive inputs—that is, the Morishima elasticity of substitution between emissions and any input $j \in \{K, L, M\}$ should be identical. We now present evidence validating this assumption.

Morishima Elasticities. Table A3 in Appendix C.4 reports estimates of the three Morishima elasticities M_{jE} for $j \in \{K, L, M\}$, estimated from regressions of $\log(j/E)$ on $\log P^E$ as described in Section 4.3. Both OLS and IV estimates are uniformly low—below 0.1 for all input pairs—and statistically indistinguishable from each other. The capital-emissions elasticity is 0.025 (IV), the labor-emissions elasticity is 0.024 (IV), and the materials-emissions elasticity is 0.015 (IV). The near-equality of these estimates supports the nested CES restriction: emissions appear to be similarly complementary with all productive inputs.

The uniformity of these estimates has a direct economic interpretation beyond validating our functional form. The consistently low Morishima elasticities imply that firms cannot easily reduce emissions by shifting toward cleaner input mixes—substituting capital for materials, for example, provides no emissions benefit. This finding reinforces the message from Section 5.1: firm-level abatement through input substitution is not a viable margin, regardless of which input pair we examine.

We also test whether the elasticity is constant across the range of observed prices by including a quadratic term $(\log P^E)^2$ in our Morishima regressions. The coefficients on the quadratic term are small and statistically insignificant (Table A4 in Appendix C.4), indicating that we cannot reject the constant elasticity specification implied by CES. Figure A6 in the appendix further shows that this low elasticity result holds across all major emitting sectors, confirming that the near-Leontief relationship between emissions and other inputs is a general feature of production in regulated industries.

Time Trend Specification. Our baseline estimation uses cubic sector-specific time trends to flexibly absorb clean technology evolution. Figure A5 in Appendix A shows that our emissions elasticity estimates are robust to alternative time trend specifications. We estimate the model with linear, quadratic, cubic, and natural spline (5 degrees of freedom) time trends. The near-zero elasticity estimates are stable across all specifications—including the highly flexible spline—indicating that our main finding of limited firm-level substitution between emissions and other inputs is not driven by the particular functional form assumed for technology trends.

6 Counterfactuals

Having estimated our model we now use it to simulate the effects of carbon pricing on the economy. Our model allows us to understand how firms and workers are affected, and what equilibrium mechanisms caused emissions reductions. In our simulations we focus on the set of firms regulated under EU-ETS at its inception in 2005 and consider three scenarios. The first is a doubling of the EU-ETS allowance price, holding clean technology fixed at 2005 levels. This doubling of the price is effectively the change in the price from 2005 to the end of our sample in 2021. The second is an increase in clean technology from 2005 to 2021 levels, but holding the emissions allowance price fixed at 2005 levels. The third is the combination of the prior two and captures both the actual increase in the allowance price and the actual improvement in clean technology.

6.1 Main Results

Table 1: Outcomes under Carbon Pricing and Clean Technology Counterfactuals

Scenario	Δ Price (%)	Δ Output (%)	Δ Emissions (%)	Δ Labor (%)
Double CO_2 Price	2.14	-2.10	-36.69	-1.62
Double CO_2 Price + Clean Tech	0.83	-0.82	-46.26	-0.70
– No Within-Industry Reallocation	1.18	-1.17	-40.14	-0.89
– No Reallocation	1.03	-1.02	-38.89	-0.87

Note: This table reports counterfactual outcomes for the set of firms regulated under EU-ETS in 2005. We simulate two scenarios: (1) a doubling of the carbon price from its 2005 level, and (2) the joint effect of doubling the carbon price and increasing clean technology to 2021 levels. We decompose the second scenario by shutting down reallocation channels: "No Within-Industry Reallocation" holds firm market shares fixed within each industry, while "No Reallocation" additionally holds industry shares fixed. All outcomes are reported as percent changes relative to baseline.

Table 1 presents aggregate results. The first row shows that doubling the carbon price leads to a 2.1 percent increase in output prices, reflecting the relatively small share of carbon in production costs. Output declines by 2.1 percent and labor demand falls by 1.6 percent. Despite these modest impacts on output and employment, emissions fall sharply by 36.7 percent. This substantial reduction is driven primarily by reallocation across firms rather than within-firm substitution—since estimated substitution elasticities are low, the technique effect accounts for only a modest share of the decline, while reallocation toward carbon-efficient producers drives the remainder.

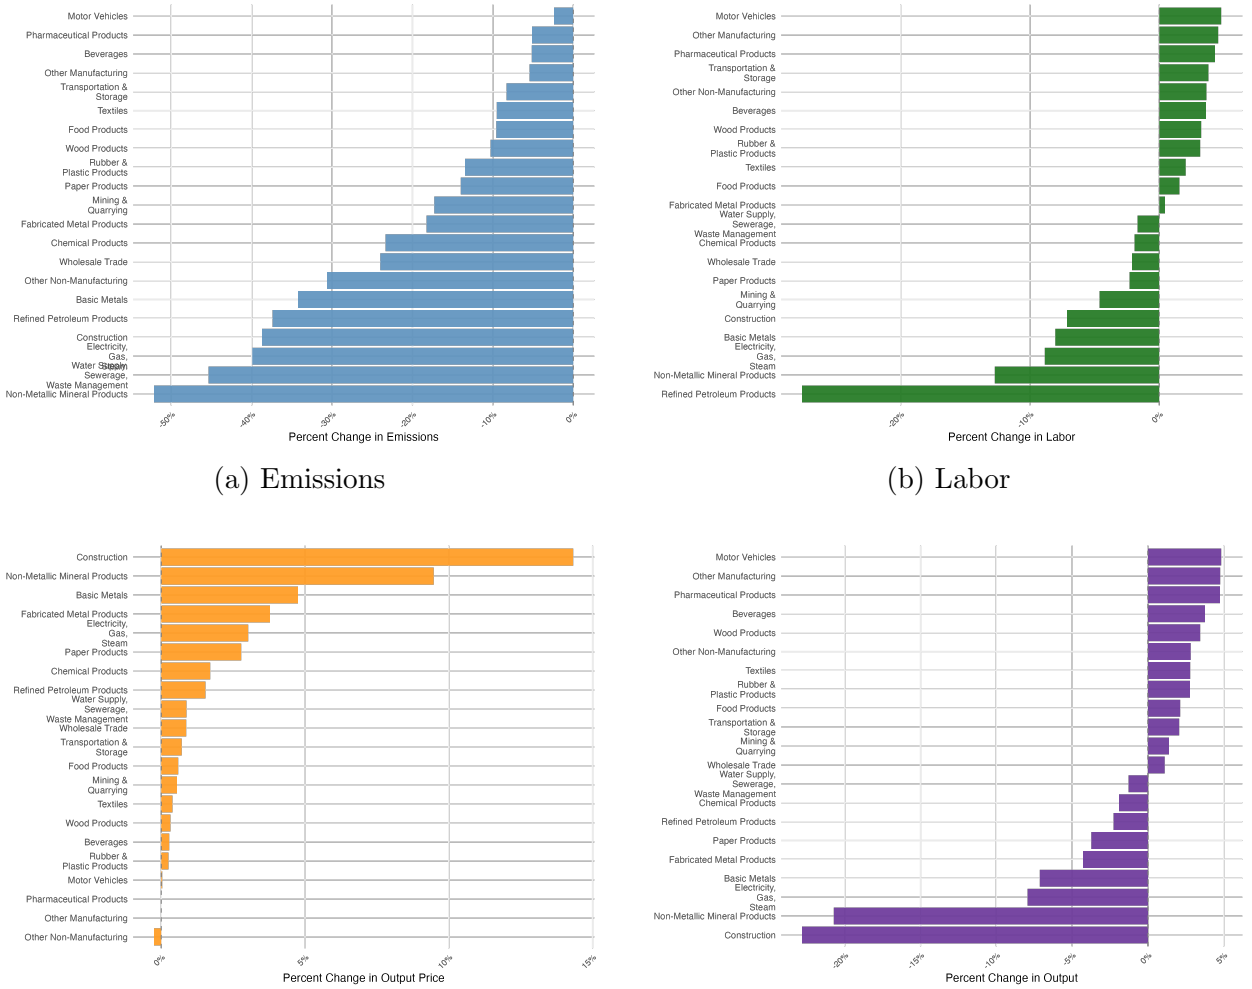
The second row shows the joint effect of doubling the carbon price and increasing clean technology to 2021 levels. Under this scenario, output prices rise by only 0.8 percent—less than half the increase under carbon pricing alone—reflecting the cost savings from cleaner

technology. Output falls by just 0.8 percent and labor demand declines by 0.7 percent. Emissions decline substantially by 46.3 percent. These results show that clean technology not only delivers larger emissions reductions but also substantially mitigates the adverse impacts of carbon pricing on firms and workers.

To understand the sources of these emissions reductions, we decompose the combined scenario by shutting down reallocation channels. Without within-industry reallocation—holding each firm's market share fixed within its industry—emissions still decline by 40.1 percent. Eliminating both within-industry and across-industry reallocation reduces emissions by 38.9 percent. The difference between the full scenario (46.3 percent) and no within-industry reallocation (40.1 percent) represents the contribution of within-industry reallocation: approximately 6 percentage points of additional emissions reduction. The difference between no within-industry reallocation (40.1 percent) and no reallocation (38.9 percent) captures across-industry reallocation: approximately 1 percentage point. These results show that within-industry reallocation—the shift in market share toward lower-emissions firms within each sector—accounts for most of the reallocation-driven emissions reduction.

6.2 Sector-Level Results

Figure 6: Sectoral Impacts of Doubling Carbon Price



(a) Emissions

(b) Labor

(c) Output Prices

(d) Output Quantities

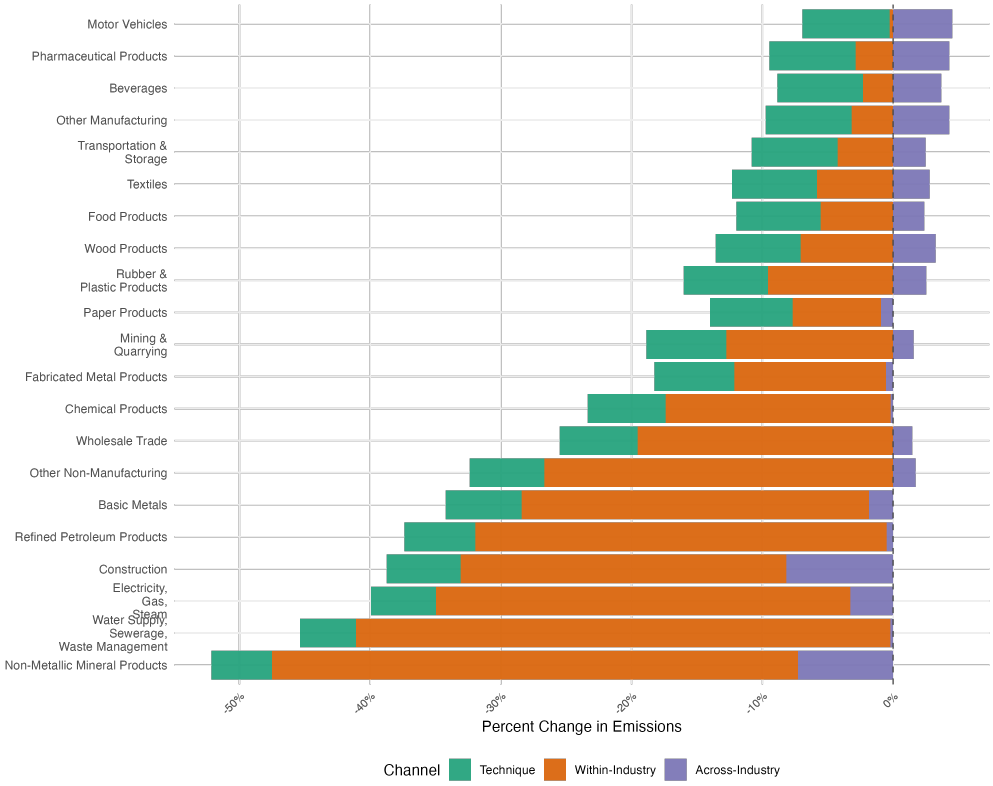
Notes: This figure plots the change in emissions (top left), labor employment (top right), output prices (bottom left), and output quantities (bottom right) by sector. In all panels, allowance prices are doubled while clean technology is held constant at 2005 levels.

Figure 6 plots the sectoral impacts of doubling the carbon price while holding clean technology constant. Across all outcomes there is substantial sectoral heterogeneity. The top left panel plots the change in emissions. All industries experience emissions reductions, but the magnitude varies considerably across sectors. Industries with the largest emissions reductions—such as refined petroleum and basic metals—tend to exhibit greater firm-level heterogeneity in emissions intensity and face more elastic demand. In contrast, sectors like electricity and gas experience smaller emissions reductions. This pattern is consistent with the role of within-industry reallocation: in heterogeneous sectors with elastic demand, carbon

pricing shifts market share toward cleaner firms, amplifying aggregate emissions reductions beyond what technique alone would deliver.

6.3 Decomposition: Role of Reallocation

Figure 7: Decomposition of Emissions Changes by Sector



Notes: This figure plots the change in emissions by mechanism. Technique captures within-firm reallocation of inputs and improvements in clean technology. Within-industry reallocation captures shifting of production from higher to lower emissions intensity firms within an industry. Across-industry reallocation captures emissions reductions from shifting consumption from higher to lower emissions intensity industry products. In the results shown in the figure, allowance prices are doubled and clean technology is held fixed.

Figure 7 decomposes emissions reductions into three different channels. The first is the technique channel which reflects how firms change their production processes—both shifting inputs into abatement activity and adopting cleaner technology. The second is within-industry reallocation which reflects how shifting production from dirtier to cleaner firms within an industry can reduce emissions even holding the distribution of emissions intensities across firms fixed. The third is across-industry reallocation which captures how emissions change as consumers substitute from industries that are dirtier on average to cleaner ones.

Within-industry reallocation is the primary driver of emissions reductions in most sectors, consistent with the aggregate decomposition in Table 1. The contribution of reallocation varies substantially across industries, with heterogeneous sectors like refined petroleum experiencing the largest reallocation effects. The technique channel—which includes both within-firm input substitution and clean technology adoption—contributes meaningfully to emissions reductions. However, within-firm substitution alone is limited by the steep firm-level abatement curves implied by our low estimated elasticities of substitution. The effect of across-industry reallocation is more muted and heterogeneous—in some cases it is positive because cleaner sectors grow so much in scale that their overall emissions increase despite lower emissions intensities.

7 Conclusion

This paper develops a framework for constructing marginal abatement cost curves from firm-level data on emissions and financials. The key methodological insight is that firm-level production parameters can be aggregated to industry and economy-wide abatement costs in a way that accounts for equilibrium reallocation across heterogeneous firms. The framework is portable: it can be applied wherever emissions pricing programs generate the requisite data, providing a template for evaluating carbon pricing policies across regulatory contexts.

Applying the framework to the EU-ETS, we find that firm-level abatement costs are steep—firms cannot easily reduce emissions without cutting output. But industry-level abatement costs are far flatter because carbon pricing shifts production toward cleaner firms. This reallocation channel operates most strongly in high-emitting sectors like electricity, petroleum, and basic metals, precisely where it matters most for climate policy. We also find that clean technology improved substantially over our sample period, shifting abatement cost curves downward and further reducing the economic costs of emissions reductions.

These findings help reconcile a puzzle: how have EU-ETS emissions fallen so sharply while output and employment remained relatively stable? The answer is that reallocation across firms and improvements in clean technology are the primary adjustment margins. Steep firm-level abatement costs coexist with modest aggregate policy costs because industries can reduce emissions far more cheaply than individual firms. For policymakers, these results suggest that carbon pricing can deliver substantial emissions reductions without severe economic disruption, especially when technological progress reduces firms' emissions intensity over time.

References

Bjørnland, Hilde Christiane, Jamie Cross, and Felix Kapfhammer (2024) *The drivers of emission reductions in the European carbon market*, No. 1/2024: Working Paper.

Blackorby, Charles and R Robert Russell (1989) "Will the real elasticity of substitution please stand up?(A comparison of the Allen/Uzawa and Morishima elasticities)," *The American economic review*, Vol. 79, No. 4, pp. 882–888.

Calel, Raphael (2020) "Adopt or innovate: Understanding technological responses to cap-and-trade," *American Economic Journal: Economic Policy*, Vol. 12, No. 3, pp. 170–201.

Calel, Raphael and Antoine Dechezleprêtre (2016) "Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market," *The Review of Economics and Statistics*, Vol. 98, No. 1, pp. 173–191.

Colmer, Jonathan, Ralf Martin, Mirabelle Muûls, and Ulrich J. Wagner (2024) "Does Pricing Carbon Mitigate Climate Change? Firm-Level Evidence from the European Union Emissions Trading System," *The Review of Economic Studies*.

Copeland, Brian R. and M. Scott Taylor (2004) "Trade, Growth, and the Environment," *Journal of Economic Literature*, Vol. 42, No. 1, pp. 7–71.

Dechezleprêtre, Antoine, David Hemous, and Morten Olsen (2023) "The Joint Impact of the European Union Emissions Trading System on Carbon Emissions and Economic Performance," *Journal of Environmental Economics and Management*, Vol. 118, p. 102572.

Gabaix, Xavier and Ralph SJ Koijen (2024) "Granular instrumental variables," *Journal of Political Economy*, Vol. 132, No. 7, pp. 2274–2303.

Hollingsworth, Alex, Taylor Jaworski, Carl Kitchens, and Ivan Rudik (2024) "Economic Geography and Air Pollution Regulation in the United States."

Jezequel-Royer, Marie and Gregory Levieuge (2025) "Going Green: Will the R-Star Shine Brighter?," https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5050773, SSRN Working Paper.

Känzig, Diego R. (2025) "The Unequal Economic Consequences of Carbon Pricing," *NBER Working Paper No. 31221*.

Letout, Simon (2021) "Firm level data in the EU ETS (JRC-EU ETS-FIRMS)," http://data.europa.eu/89h/bdd1b71f-1bc8-4e65-8123-bbdd8981f116, European Commission, Joint Research Centre (JRC) [Dataset].

Marin, Giovanni, Marianna Marino, and Claudia Pellegrin (2018) "The impact of the European Emission Trading Scheme on multiple measures of economic performance," *Environmental and Resource Economics*, Vol. 71, pp. 551–582.

Metcalf, Gilbert E. and James H. Stock (2023) "The Macroeconomic Impact of Europe's Carbon Taxes," *American Economic Journal: Macroeconomics*, Vol. 15, No. 3, pp. 265–286.

Nakamura, Emi and Jón Steinsson (2018) "Identification in macroeconomics," *Journal of Economic Perspectives*, Vol. 32, No. 3, pp. 59–86.

Oberfield, Ezra and Devesh Raval (2021) "Micro Data and Macro Technology," *Econometrica*, Vol. 89, No. 2, pp. 703–732.

Ortubai, Joxe Mari Barrutiabengoa, Rodrigo Enrique Falbo Piacentini, Agustın Garcıa Serrador, and Juan F Rubio-Ramırez (2025) "Unraveling the impact of a carbon price shock on macroeconomic variables: a Narrative Sign Restrictions approach."."

Petrick, Sebastian and Ulrich J Wagner (2014) "The impact of carbon trading on industry: Evidence from German manufacturing firms," *Available at SSRN 2389800*.

Raval, Devesh R. (2019) "The micro elasticity of substitution and non-neutral technology," *The RAND Journal of Economics*, Vol. 50, No. 1, pp. 147–167.

Shapiro, Joseph S. and Reed Walker (2018) "Why Is Pollution from U.S. Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade," *American Economic Review*, Vol. 108, No. 12, pp. 3814–3854.

Appendix Table of Contents

A Additional Tables and Figures

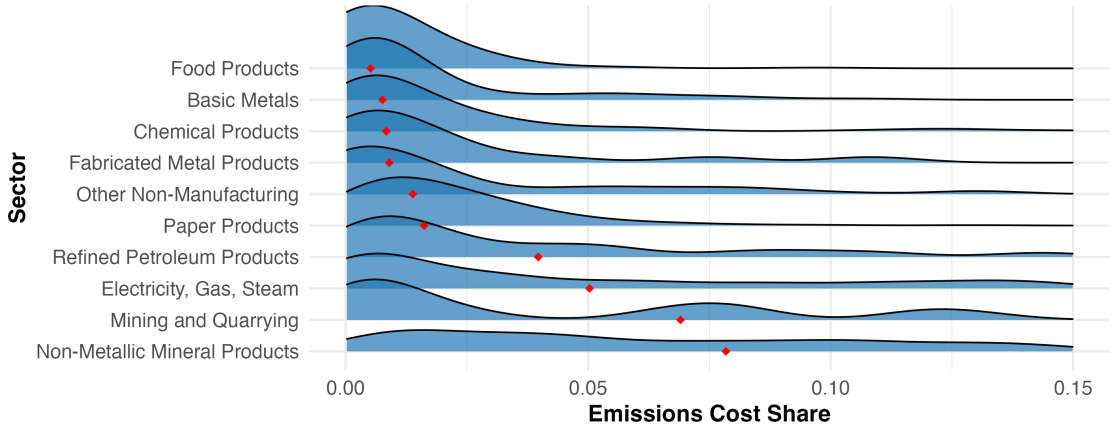
B Statistical Decomposition of Emissions

C Generalized Model of Production and Abatement

?? Aggregation and Marginal Abatement Cost Curves

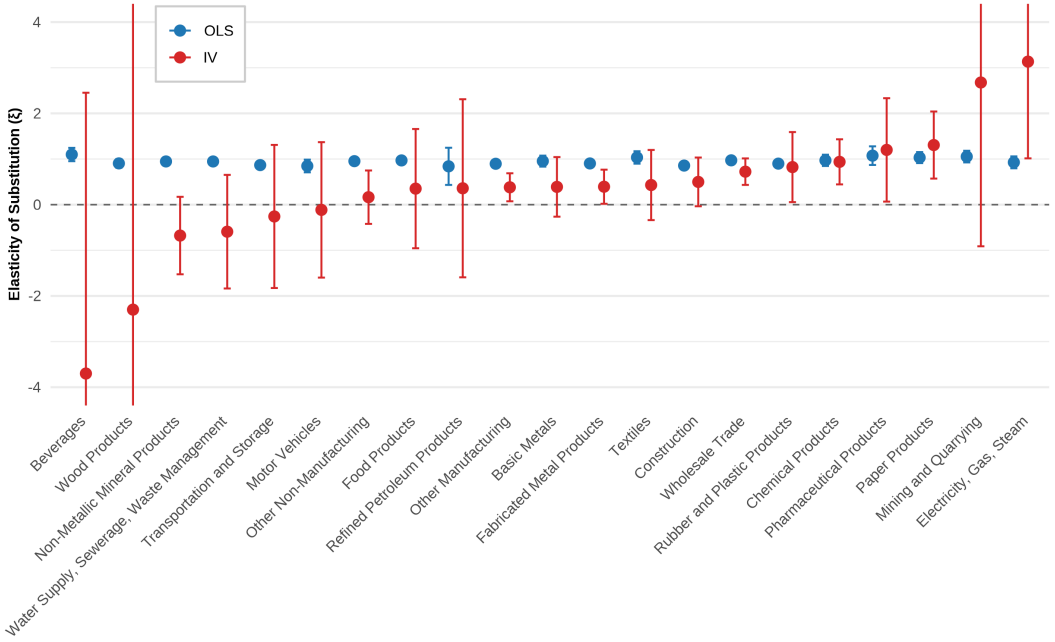
A Additional Tables and Figures

Figure A1: Distribution of Emissions Cost Shares Across Industries



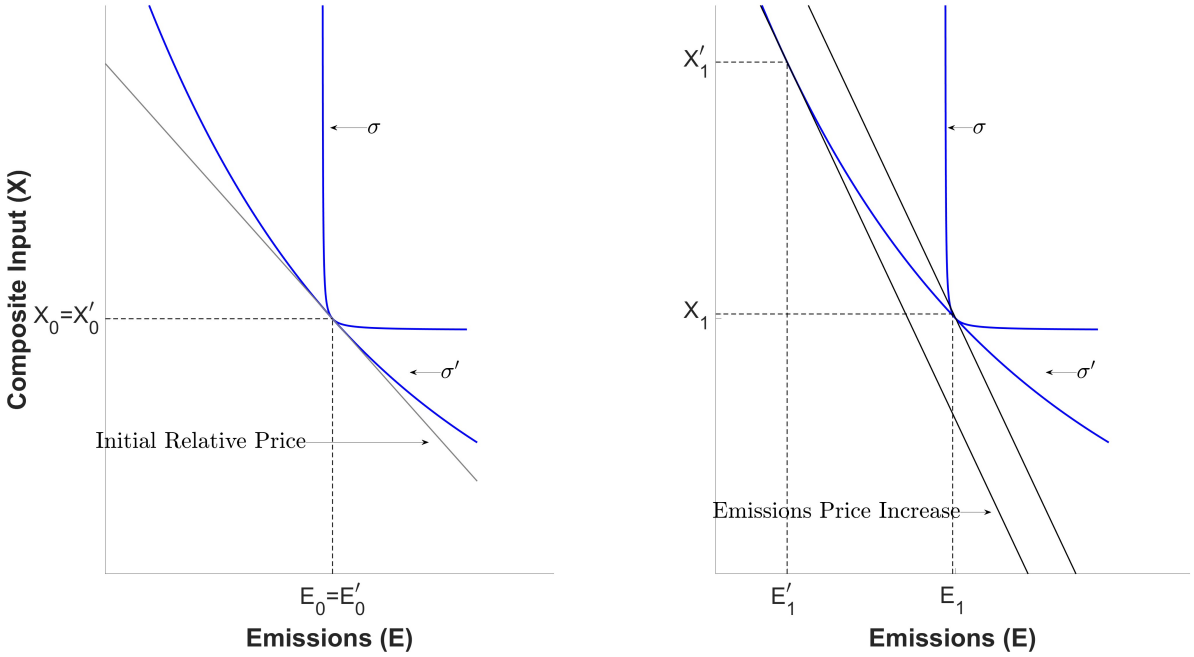
Notes: This figure shows the distribution of emissions cost shares across firms within the ten highest-emitting industries in 2005. Each ridge depicts the within-industry distribution of the ratio of emissions cost (tons of CO_2 times the average permit transaction price) to firms' cost of goods sold. Red dots mark the median firm. In sectors such as petroleum products or non-metallic minerals, the distributions are wide and right-skewed, indicating substantial heterogeneity in emissions intensity and scope for emissions reductions through output reallocation from dirtier to cleaner firms. In contrast, more homogeneous sectors like food products offer less opportunity for such reallocation-based abatement.

Figure A2: Capital-Labor-Materials Elasticity of Substitution by Sector



Notes: This figure plots sector-specific estimates of the inner-nest elasticity of substitution ξ_n among capital, labor, and materials from equation (20). Blue circles show OLS estimates; red diamonds show IV estimates using granular instrumental variables for wages following Gabaix and Koijen (2024). Error bars represent 95% confidence intervals based on standard errors clustered at the firm level. The dashed vertical line marks one (Cobb-Douglas).

Figure A3: Elasticity of Substitution Between Emissions and Composite Input



(a) Initial Relative Price

(b) Emissions Price Increase

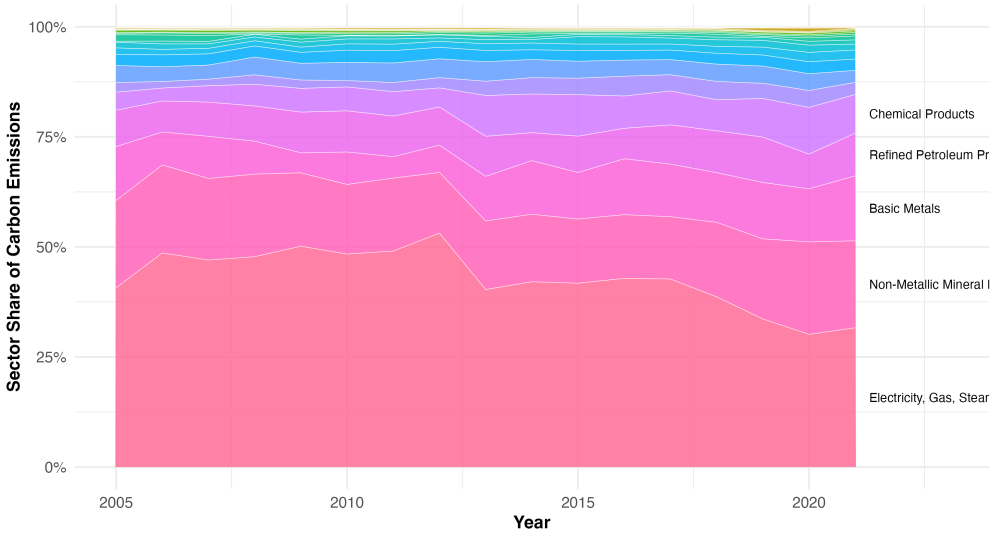
Notes: This figure illustrates input responses to an emissions price increase for firms with different production technologies. The isoquant labeled σ has low substitutability between the composite input (X) and emissions (E); the isoquant labeled σ' has higher substitutability. Panel (a) shows optimal input choices (X_0, E_0 and X'_0, E'_0) under an initial relative price (gray line). Panel (b) shows responses (X_1, E_1 and X'_1, E'_1) to an emissions price increase (black line). With low σ, emissions barely change; with higher σ', firms substantially reduce emissions by substituting toward the composite input.

Table A1: Sector Summary

Sector	NACE Codes	# Observations	# Unique Firms	Share of Emissions (%)
Electricity, Gas, Steam	35	8,603	1,004	42.95
Non-Metallic Mineral Products	23	8,435	1,002	16.97
Paper Products	17	4,196	443	2.5
Chemical Products	20	3,122	348	6.7
Food Products	10	3,036	338	1.49
Basic Metals	24	2,112	242	9.94
Other Non-Manufacturing	01–02, 45, 47, 61–64, 68–74, 77–82, 85–86, 90, 93, 96	1,228	206	2.86
Other Manufacturing	12, 14–15, 18, 26–28, 30–33	1,137	159	0.33
Wholesale Trade	46	754	133	1.24
Wood Products	16	705	95	0.17
Mining and Quarrying	05–09	658	64	3.82
Textiles	13	585	76	0.06
Water Supply, Sewerage, Waste Management	36–39	530	67	0.23
Refined Petroleum Products	19	529	55	8.5
Fabricated Metal Products	25	485	66	0.59
Rubber and Plastic Products	22	483	58	0.17
Beverages	11	478	52	0.09
Motor Vehicles	29	448	47	0.26
Pharmaceutical Products	21	428	45	0.13
Transportation and Storage	49, 52	420	53	0.38
Construction	41–43	414	68	0.62

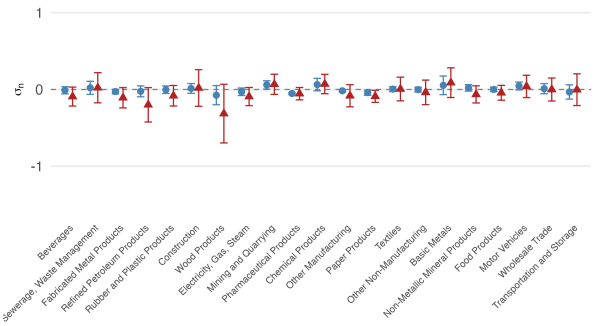
Notes: This table summarizes sectoral coverage in the matched EU-ETS–Orbis sample, reporting the number of firm-year observations, unique firms, and each sector's share of total emissions. Emissions shares are computed over the full sample period.

Figure A4: Emissions Share by Sector Over Time

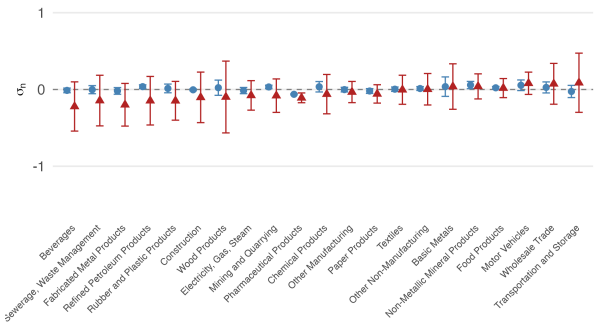


Notes: This figure plots each sector's share of total verified emissions under EU-ETS over time.

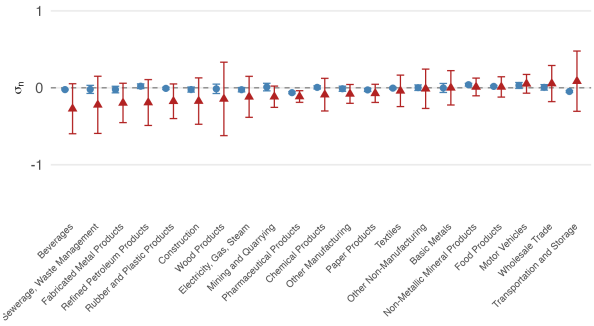
Figure A5: Emissions Elasticity Robustness to Time Trend Specification



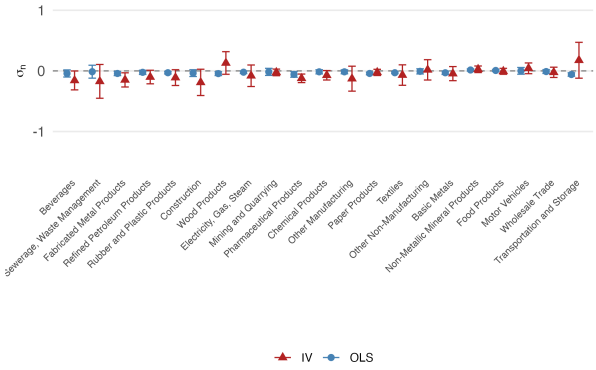
(a) Linear Time Trend



(b) Quadratic Time Trend



(c) Cubic Time Trend (Baseline)



(d) Natural Spline (5 df)

Notes: This figure shows emissions-composite elasticity estimates under alternative time trend specifications. Blue circles show OLS estimates; red triangles show IV estimates using the Jezequel-Royer and Levieuge (2025) carbon price surprise instrument. Error bars represent 95% confidence intervals based on standard errors clustered at the firm and year level. The estimates are stable across specifications, with all sectors showing near-zero elasticities regardless of the time trend flexibility.

Table A2: Cross-Industry Demand Elasticity Estimates

	Log(Q)			
	(1)	(2)	(3)	(4)
Log(P)	-0.4130	-0.4276	-2.764	-3.221
	(0.3452)	(0.3731)	(0.5788)	(0.4483)
Model	OLS	OLS	IV	IV
NACE Section by Year Trends	No	Yes	No	Yes
Observations	266	266	266	266
NACE 2-Digit fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes

Notes: This table reports estimates of the cross-industry demand elasticity υ from the regression $\log Q_{nt} = -\upsilon \log P_{nt} + h_n(t) + \varepsilon_{nt}$, where Q_{nt} and P_{nt} are sectoral quantity and price indices from EU KLEMS. Columns (1)–(2) report OLS estimates; columns (3)–(4) report IV estimates using sector average labor productivity as an instrument for price. Columns (2) and (4) include NACE section-by-year trends. All specifications include NACE 2-digit sector fixed effects and year fixed effects with standard errors clustered at the NACE 2-digit level. Column (4) is the baseline specification used in counterfactual simulations. Standard errors in parentheses.

B Statistical Decomposition of Emissions

C Generalized Model of Production and Abatement

The main text adopts a nested CES production structure in which emissions enter the outer nest alongside a composite of capital, labor, and materials. This functional form imposes two key restrictions: (i) a constant elasticity of substitution between emissions and other inputs, and (ii) identical substitutability between emissions and each productive input. This appendix provides theoretical and empirical foundations for these choices.

We first present a general production model that places minimal structure on the production process, requiring only constant returns to scale and a particular—but still fairly general—relationship between abatement effort and emissions. We prove that any such model is isomorphic to a production function with emissions as an input, which means our estimation approach can be applied to a broad class of production specifications beyond

CES.

While the isomorphism result justifies treating emissions as an input, it does not pin down the functional form. Importantly, the Morishima elasticity of substitution—the appropriate measure of pairwise substitutability with more than two inputs—can be estimated without imposing any functional form on production. However, to simulate counterfactuals we must impose a specific functional form. The nested CES is a natural candidate: it assumes that emissions are equally substitutable with all productive inputs (capital, labor, and materials) and that this substitutability is constant. We test these restrictions by estimating Morishima elasticities between emissions and each productive input. We show empirically that these elasticities are uniformly low, approximately constant across price levels, and statistically indistinguishable across input pairs. Hence, the Morishima elasticity estimates provide direct empirical support for our choice of the nested CES specification in the main text.

Section C.1 introduces the generalized production and abatement framework. Section C.2 states the key isomorphism result showing equivalence to a production function with emissions as an input. Section C.3 provides the formal proof. Section C.4 presents empirical tests validating our nested CES specification.

C.1 Production and Emissions Abatement

Sectors are indexed by $n \in 1, \ldots, N$ and firms within a sector n are indexed by $i \in 1, \ldots, F_n$. The timestep in our model is a year and years are indexed by t. A firm uses capital K_{int}, labor L_{int}, and materials M_{int}, allocating them to production of output or abatement of emissions which are a byproduct of production:

$$K_{int} = K_{y,int} + K_{a,int}, \qquad L_{int} = L_{y,int} + L_{a,int}, \qquad M_{int} = M_{y,int} + M_{a,int}$$

where all allocations are nonnegative. We denote the vector of total inputs by $v_{int} = (K_{int}, L_{int}, M_{int})$. Let X be a continuous, increasing, concave, and constant returns to scale production function:

$$Y_{int} \;=\; A_{int}\, X(K_{y,int}, L_{y,int}, M_{y,int}),$$

where $A_{int} > 0$ is Hicks-neutral productivity. Let Z be an abatement function. Z is continuous, increasing, and constant returns to scale for each B_{int}, and increasing in the clean technology B_{int}. Abatement is then:

$$Z_{a,int} \equiv Z(K_{a,int}, L_{a,int}, M_{a,int}; B_{int}).$$

Better clean technology B_{int} will generate more abatement given the same amount of inputs allocated toward abatement.

We next impose a relationship between emissions, production, and abatement. Emissions are proportional to potential output and depend on a function of abatement intensity:[10]

$$E_{int} \;=\; A_{int}\, X(v_{int})\, \psi\!\left(\frac{A_{int}\, Z_{a,int}}{Y_{int}}\right) \tag{32}$$

The argument $A_{int}Z_{a,int}/Y_{int}$ is the firm's abatement intensity: the ratio of productivity-adjusted abatement to actual output. Scaling abatement by A_{int} ensures that the argument is homogeneous of degree zero in the firm's scale, so that abatement intensity depends only on how the firm allocates its inputs between production and abatement, not on its overall size.

Finally, we impose a few assumptions on the properties of abatement function ψ:

Assumption 1 *ψ ranges between $(\underline{e}, \overline{e})$ and is continuous, strictly decreasing, strictly convex, with $0 \leq \underline{e} < \overline{e} < \infty$.*

Decreasing ψ means emissions decrease in the amount of abatement relative to production. Strict convexity imposes increasing marginal abatement costs. The bounds ensure emissions are weakly positive and that emissions without any abatement is finite.

C.2 Production Isomorphism

The framework above places minimal structure on how firms produce output and abate emissions: it requires only that production and abatement technologies exhibit constant returns to scale and that emissions depend on abatement intensity as specified in Assumption 1. A key challenge in taking this model to the data is that we do not observe how firms divide their inputs between production and abatement. We show that this is not a problem: the setup is isomorphic to one in which emissions are treated as an input to a concave, constant returns to scale production function. This equivalence allows us to estimate the model using observed total inputs and emissions, without needing to observe the underlying allocation.

The key steps of the proof are: (1) derive the cost function dual to the original production setup for achieving some combination of output and emissions; (2) define G as the maximal feasible output given some combination of inputs and emissions; (3) show that G is increasing in emissions; (4) show that the cost function dual to G equals the original cost function everywhere. The formal proof is provided in Section C.3.

[10] $A_{int}\, X(v_{int})$ is potential output since it is how much is produced if all inputs were allocated toward production instead of abatement.

This isomorphism result is powerful: it allows us to estimate the supply side of the model using standard production function techniques applied to observed $(K_{int}, L_{int}, M_{int}, E_{int})$, without needing to observe abatement directly or how firms allocate inputs between production and abatement activities.

C.3 Proof of Production Isomorphism

This section provides a formal proof that the general production and abatement framework is isomorphic to a production function with emissions as an input. We first derive the cost function for achieving some combination of output and emissions, then prove that its properties imply there exists a production function G with emissions as an input. Recall that $v_{int} = (K_{int}, L_{int}, M_{int})$ denotes the vector of total inputs.

Deriving the Cost Function

Because ψ is strictly monotone, it is invertible and its inverse has domain $(\underline{e}, \overline{e})$. Let $\phi \equiv \psi^{-1}$. ϕ is strictly decreasing and strictly convex since ψ is strictly decreasing and strictly convex. From equation (32) we can back out the abatement requirement to achieve any given output-emissions (Y_{int}, E_{int}) combination:

$$Z_{a,int} \;=\; \frac{Y_{int}}{A_{int}}\,\phi\Big(\frac{A_{int}E_{int}}{Y_{int}}\Big). \tag{33}$$

The ratio $e_{int} \equiv A_{int}E_{int}/Y_{int}$ is emissions per unit of the productive composite X. Next we derive the minimum cost of achieving some (Y_{int}, E_{int}). Let $p_{int} = (P^K_{int}, P^L_{int}, P^M_{int})$ be a vector of input prices. Let $c_X(p_{int})$ and $c_Z(p_{int}; B_{int})$ be the unit cost functions for production and abatement. The unit cost functions are homogeneous degree one and concave:

$$c_X(p_{int}) = \inf_{u\geq 0}\{p_{int}\cdot u : X(u)\geq 1\} \qquad c_Z(p_{int}; B_{int}) = \inf_{w\geq 0}\{p_{int}\cdot w : Z(w; B_{int})\geq 1\} \tag{34}$$

where c_Z is also decreasing in B_{int} since X is increasing in B_{int}. Because production and abatement are separable, the minimum total cost of delivering (Y_{int}, E_{int}) is the minimum cost of producing Y_{int}/A_{int} plus the minimum cost of producing $Z_{a,int}$ from equation (33):

$$C(Y_{int}, E_{int};\, p_{int}, A_{int}, B_{int}) \;=\; c_X(p_{int})\,\frac{Y_{int}}{A_{int}} + c_Z(p_{int}; B_{int})\,Z_{a,int} \;=\; \frac{Y_{int}}{A_{int}}c_X(p_{int}) + c_Z(p_{int}; B_{int})\,\frac{Y_{int}}{A_{int}}\,\phi\Big(\frac{A} \tag{35}$$

Let g be defined as:

$$g(Y_{int}, E_{int}; A_{int}) \equiv \frac{Y_{int}}{A_{int}} \phi\left(\frac{A_{int}E_{int}}{Y_{int}}\right).$$

g converts the per-unit-production abatement requirement $\phi(A_{int}E_{int}/Y_{int})$ into a total requirement by scaling with production. g will be convex since ϕ is convex. We show this in Lemma 1:

Lemma 1 *If ϕ is convex, then $g(Y_{int}, E_{int}; A_{int})$ is jointly convex in (Y_{int}, E_{int}).*

Proof: Let $(Y_1, E_1), (Y_2, E_2)$ and $\lambda \in [0, 1]$. Let $Y_\lambda = \lambda Y_1 + (1-\lambda)Y_2 > 0$ and let $E_\lambda = \lambda E_1 + (1-\lambda)E_2$. Next, define convex combination weights $w_1 \equiv \lambda Y_1/Y_\lambda$, $w_2 \equiv (1-\lambda)Y_2/Y_\lambda$. Clearly $w_1, w_2 \geq 0$ and $w_1 + w_2 = 1$, and we also have that:

$$\frac{AE_\lambda}{Y_\lambda} = w_1 \frac{AE_1}{Y_1} + w_2 \frac{AE_2}{Y_2}.$$

Since ϕ is convex we have that:

$$\phi\big(AE_\lambda/Y_\lambda\big) \leq w_1\, \phi(AE_1/Y_1) + w_2\, \phi(AE_2/Y_2).$$

Multiply both sides by Y_λ/A:

$$\underbrace{\frac{Y_\lambda}{A}\phi\big(AE_\lambda/Y_\lambda\big)}_{g(Y_\lambda, E_\lambda)} \leq \underbrace{\frac{Y_\lambda}{A}w_1}_{\lambda\frac{Y_1}{A}}\ \phi(AE_1/Y_1) + \underbrace{\frac{Y_\lambda}{A}w_2}_{(1-\lambda)\frac{Y_2}{A}}\ \phi(AE_2/Y_2) \tag{36}$$

which then reduces down to the definition of convexity for g:

$$g(Y_\lambda, E_\lambda) \leq \lambda g(Y_1, E_1) + (1-\lambda) g(Y_2, E_2).$$

□

With convexity of g, we now introduce Proposition 1 which shows the key properties of the cost function that will deliver the emissions-as-an-input production function representation:

Proposition 1 *The cost function in equation* (35) *has the following five properties:*

1. *Convexity in (Y_{int}, E_{int}): $C(\cdot, \cdot;\, p_{int}, A_{int}, B_{int})$ is convex.*
2. *Constant returns to scale in (Y_{int}, E_{int}): For any $t > 0$, $C(tY_{int}, tE_{int};\, p_{int}, A_{int}, B_{int}) = t\, C(Y_{int}, E_{int};\, p_{int}, A_{int}, B_{int})$.*

3. *Decreasing in* E_{int}*:* $\dfrac{\partial C}{\partial E_{int}}(Y_{int}, E_{int};\, p_{int}, A_{int}, B_{int}) = c_Z(p_{int}; B_{int})\,\phi'\!\left(\frac{A_{int}E_{int}}{Y_{int}}\right) < 0.$

4. C *is nondecreasing in each price component of* p_{int} *and weakly decreasing in* B_{int}*.*

5. *Boundary at* $Y_{int} = 0$*: Extend* g *by lower-semicontinuous closure via* $g(0,0) = 0$ *and* $g(0, E_{int} > 0) = +\infty$*; then* $C(0, 0;\, p_{int}, A_{int}, B_{int}) = 0$ *and* $C(0, E_{int} > 0;\, p_{int}, A_{int}, B_{int}) = +\infty$*.*

Proof:

1. g's convexity was established in Lemma 1. Since $C(Y_{int}, E_{int};\, p_{int}, A_{int}, B_{int}) = c_X(p_{int})\, Y_{int}/A_{int} + c_Z(p_{int}; B_{int})\, g(Y_{int}, E_{int}; A_{int})$, C is the sum of a linear term and a nonnegative multiple of a convex function so C is convex in (Y_{int}, E_{int}).

2. Multiply Y_{int} and E_{int} by t in equation (35) to obtain:

$$C(tY_{int}, tE_{int};\, p_{int}, A_{int}, B_{int}) = \frac{tY_{int}}{A_{int}}\left[c_X(p_{int}) + c_Z(p_{int}; B_{int})\,\phi\!\left(\frac{A_{int}(tE_{int})}{tY_{int}}\right)\right] = t\,C(Y_{int}, E_{int};\, p_{int}, A_{i}$$

3. Differentiate equation (35) with respect to E_{int} holding prices fixed:

$$\frac{\partial C}{\partial E_{int}}(Y_{int}, E_{int};\, p_{int}, A_{int}, B_{int}) = \frac{Y_{int}}{A_{int}}\, c_Z(p_{int}; B_{int})\,\phi'\!\left(\frac{A_{int}E_{int}}{Y_{int}}\right)\!\cdot\frac{A_{int}}{Y_{int}} = c_Z(p_{int}; B_{int})\,\phi'\!\left(\frac{A_{int}E_{int}}{Y_{int}}\right) <$$

since $c_Z > 0$ and $\phi' < 0$.

4. Since c_X and c_Z are unit-cost functions they are nondecreasing in prices. C is then nondecreasing in each element of p_{int} since Y_{int}/A_{int} and $\phi(\cdot)$ are both nonnegative making it the sum of two nondecreasing functions of prices. Since $c_Z(p_{int}; B_{int})$ is weakly decreasing in B_{int}, C is weakly decreasing in B_{int} since ϕ is nonnegative.

5. Lower-semicontinuous extension of g at the boundary preserves convexity. With that extension, $C(0, 0;\, p_{int}, A_{int}, B_{int}) = 0$ by inspection, and $C(0, E_{int} > 0;\, p_{int}, A_{int}, B_{int}) = +\infty$ since positive emissions cannot be produced at finite cost with zero output.

□

What does each part of the proposition buy us for our emissions-as-an-input production function? Part 1 buys us that we will have convex production sets. Part 2 gives us constant returns to scale. Part 3 gives us that emissions behave like a productive input: if cost falls with emissions then that means emissions increase production, all else equal. Part 4 gives

us the clear comparative statics we need. Part 5 pins down the origin and requires positive output for positive emissions.

C.3.1 Isomorphic Production Function with Emissions as an Input

Now we show there is a production function that is isomorphic to the original production set up but where emissions is treated as an input.

Proposition 2 *Let $C(Y_{int}, E_{int};\, p_{int}, A_{int}, B_{int})$ be the cost function in (35) that inherits the properties in Proposition 1. Let $v_{int} \equiv (K_{int}, L_{int}, M_{int})$ be the vector of productive inputs. Define the feasible set:*

$$\mathcal{F} \;\equiv\; \bigcap_{p_{int}>0} \Big\{(Y_{int}, E_{int}, v_{int}) :\; C(Y_{int}, E_{int};\, p_{int}, A_{int}, B_{int}) \;\leq\; p_{int} \cdot v_{int}\Big\}$$

and define the production function with emissions as an input:

$$G(v_{int}, E_{int}; A_{int}, B_{int}) \;\equiv\; \sup\{\, Y_{int} \geq 0 :\; (Y_{int}, E_{int}, v_{int}) \in \mathcal{F}\}.$$

Then the following five properties hold:

1. *Concavity in (v_{int}, E_{int}): $G(\cdot, \cdot; A_{int}, B_{int})$ is concave.*
2. *G is constant returns to scale in (v_{int}, E_{int}): For any $t > 0$, $G(tv_{int}, tE_{int}; A_{int}, B_{int}) = t\, G(v_{int}, E_{int}; A_{int}, B_{int})$.*
3. *Output increases in emissions: $G(v_{int}, E'_{int}; A_{int}, B_{int}) \geq G(v_{int}, E_{int}; A_{int}, B_{int})$ if $E'_{int} \geq E_{int}$.*
4. *Duality with C: For all (Y_{int}, E_{int}), the cost function dual to $G(v_{int}, E_{int}; A_{int}, B_{int})$ is $C(Y_{int}, E_{int}; p_{int}, B_{int})$:*

$$\inf_{v_{int} \geq 0} \{\, p_{int} \cdot v_{int} :\; G(v_{int}, E_{int}; A_{int}, B_{int}) \geq Y_{int}\,\} = C(Y_{int}, E_{int};\, p_{int}, A_{int}, B_{int}).$$

Proof:

1. Assume input prices p are strictly positive. Since $C(\cdot, \cdot;\, p, A, B)$ is convex in (Y, E) as shown in Proposition 1.1, and thus also lower semicontinuous in (Y, E), the set:

$$\{(Y, E, v) : C(Y, E;\, p, A, B) \leq p \cdot v\}$$

is a closed half-space in (Y, E, v). Intersecting over all p yields a nonempty, closed, convex feasible set $\mathcal{F}$. Nonemptiness follows from Proposition 1.5 showing that (0,0,0) is a part of $\mathcal{F}$. The upper boundary in the Y-direction of a convex set is concave, hence G is concave in (v, E).

2. Recall that constant returns to scale of C implies $C(tY, tE;\, p, A, B) = t\, C(Y, E;\, p, A, B)$ for $t > 0$. Consider some allocation (Y, E, v) is in the feasible set so that $C(Y, E;\, p, A, B) \leq p \cdot v$ for all p. Multiply both sides by some $t > 0$. Using constant returns to scale of C we have that: $C(tY, tE;\, p, A, B) \leq p \cdot (tv)$ for all p. Thus any positive scalar multiple of (Y, E, v) is in the feasible set: $(tY, tE, tv) \in \mathcal{F}$. Therefore $\mathcal{F}$ is a cone.

   Next, define the set of feasible output levels given a level of inputs and emissions:

$$S(v, E) \;\equiv\; \{\, Y \geq 0 :\; (Y, E, v) \in \mathcal{F} \,\}.$$

   By definition, $G(v, E) = \sup S(v, E)$.

   Clearly we have that $Y \in S(v, E)$ if and only if $(Y, E, v) \in \mathcal{F}$ by the definition of S. Since $\mathcal{F}$ is a cone, we have that $Y \in S(v, E)$ if and only if $(tY, tE, tv) \in \mathcal{F}$. Finally, by defintion of S, this is equivalent to $tY \in S(tv, tE)$. Combining all the equivalences gives us that: $Y \in S(v, E) \iff tY \in S(tv, tE)$. This then yields that:

$$S(tv, tE) \;=\; t\, S(v, E)$$

   where $tS(v, E)$ is defined as:

$$t\, S(v, E) \;=\; \{\, tY :\; Y \in S(v, E) \,\}.$$

   Recognizing that under a supremum: $\sup\{tY :\; Y \in A\} = t \sup A$ for $t > 0$, we then have that:

$$G(tv, tE; A, B) \;=\; \sup S(tv, tE) \;=\; \sup\{tY :\; Y \in S(v, E)\} \;=\; t \sup S(v, E) \;=\; t\, G(v, E; A, B).$$

3. Assume $E' \geq E$. If $(Y, E, v) \in \mathcal{F}$ then by definition $C(Y, E;\, p, A, B) \leq p \cdot v$. Since C is decreasing in E as shown in Proposition 1.3, for any $E' \geq E$ we have that $C(Y, E';\, p, A, B) \leq C(Y, E;\, p, A, B) \leq p \cdot v$ so $(Y, E', v) \in \mathcal{F}$. Since $Y \in S(v, E')$ if $Y \in S(v, E)$, we then have that $S(v, E) \subseteq S(v, E')$.

Since G yields the maximal Y in the feasible set:

$$G(v, E; A, B) \equiv \sup\{ Y \geq 0 : \ (Y, E, v) \in \mathcal{F}\}.$$

it must be that $G(v, E'; A, B) \geq G(v, E; A, B)$ since the supremum of a superset cannot be smaller.

4. Define the cost function dual to G as:

$$e(Y, E;\, p, A, B) \equiv \inf_{v \geq 0}\{p \cdot v : \ G(v, E; A, B) \geq Y\}.$$

We will show this is equivalent to $C(Y, E; p, B)$.

First we show the minimum cost e under G is at least as large as the minimum cost C of the original production structure. G is the maximum feasible output by definition. So if for some (Y, E, v) we have that $G(v, E; A, B) \geq Y$, we then have that (Y, E, v) is feasible: $(Y, E, v) \in \mathcal{F}$. Then, by definition of the feasible set $\mathcal{F}$ we have that: $C(Y, E;\, q, B) \leq q \cdot v$ for any positive price vector q. Since this holds for all q, it holds for p: $C(Y, E;\, p, A, B) \leq p \cdot v$. Taking the infimum of $p \cdot v$ over v yields $e(Y, E;\, p, A, B)$, but since $C(Y, E;\, p, A, B) \leq p \cdot v$ we have that $e(Y, E;\, p, A, B) \geq C(Y, E;\, p, A, B)$.

Next we show that the minimum cost implied by G is no larger than the minimum cost C of the original production structure by constructing a feasible input bundle v with cost that is arbitrarily close to $C(Y, E;\, p, A, B)$. Recall that minimum cost of the original production structure in equation (35) is:

$$C(Y, E;\, p, A, B) \ = \ \frac{Y}{A} c_X(p) + c_Z(p; B)\, \frac{Y}{A}\, \phi\Big(\frac{AE}{Y}\Big).$$

By the definition of the unit cost as an infimum in equation (34), we have that for all ϵ, there exists bundles u_ϵ and w_ϵ such that:

$$p \cdot u_\epsilon \leq c_F(p) + \epsilon \qquad p \cdot w_\epsilon \leq c_Z(p; b) + \epsilon \tag{37}$$

so that we can find bundles arbitrarily close to the minimum unit cost bundles defined by the unit cost functions. Next, define the cost of these ϵ-bundles as v_ϵ:

$$v_\epsilon \ \equiv \ \frac{Y}{A}\, u_\epsilon \ + \ \frac{Y}{A}\, \phi\Big(\frac{AE}{Y}\Big)\, w_\epsilon.$$

Given some price vector $q > 0$, we have that:

$$\begin{aligned} q \cdot v_\epsilon &= \frac{Y}{A} q u_\epsilon \; + \; \frac{Y}{A} \phi\Big(\frac{AE}{Y}\Big) q w_\epsilon \\ &\geq \frac{Y}{A} c_F(q) \; + \; \frac{Y}{A} \phi\Big(\frac{AE}{Y}\Big) c_Z(q; B) \\ &= C(Y, E; q, A, B). \end{aligned} \tag{38}$$

Since $q \cdot v_\epsilon > C(Y, E; q, A, B)$, by the definition of $\mathcal{F}$ we have that (Y, E, v_ϵ) is in the feasible set $(Y, E, v_\epsilon) \in \mathcal{F}$ and then we also have that $G(v_\epsilon, E; A, B) \geq Y$ by definition of G. Next, evaluate the ϵ-bundles at price p:

$$\begin{aligned} p \cdot v_\epsilon &= \frac{Y}{A} p \cdot u_\epsilon \; + \; \frac{Y}{A} \phi\Big(\frac{AE}{Y}\Big) p \cdot \; w_\epsilon \\ &\leq \frac{Y}{A} c_F(q) \; + \; \frac{Y}{A} \phi\Big(\frac{AE}{Y}\Big) c_Z(q; B) + \epsilon \frac{Y}{A} \left[1 + \phi\Big(\frac{AE}{Y}\Big)\right] \end{aligned} \tag{39}$$

where the inequality is from inserting equation (37). Take the limit as $\epsilon \downarrow 0$ to get that $\lim_{\epsilon \downarrow 0} p \cdot v_\epsilon = C(Y, E; p, B)$. Since v_ϵ is feasible and $e(Y, E; p, B)$ is defined as the infimum of $p \cdot v$ across all feasible bundles v, we have that $e(Y, E; p, B) \leq C(Y, E; p, B)$. Since $e(Y, E; p, B) \leq C(Y, E; p, B)$ and $e(Y, E; p, B) \geq C(Y, E; p, B)$ we have that $e(Y, E; p, B) = C(Y, E; p, B)$.

□

Proposition 2 establishes that G is a well-behaved concave, constant returns to scale production function where emissions behave like a productive input (since G is increasing in E). The duality result shows that G is equivalent to the original production structure: they have the same cost function, so firms exhibit identical behavior under either framework. This allows us to estimate the production parameters using observed inputs and emissions without needing to observe how inputs are allocated between production and abatement.

C.4 Morishima Elasticities: Empirical Validation of Nested CES Specification

Having established that we can treat emissions as an input without observing abatement directly, we now turn to validating our choice of the nested CES functional form. The nested CES specification in the main text assumes that the elasticity of substitution between emissions and any productive input (capital, labor, or materials) is identical. We test this

assumption using Morishima elasticities of substitution.

The Morishima elasticity of substitution M_{jE} measures how the ratio of input j to emissions responds to a change in the relative price of emissions, holding output constant:

$$M_{jE} = \left. \frac{\partial \log(j/E)}{\partial \log(P^E/P_j)} \right|_Y$$

where $j \in \{K, L, M\}$. With more than two inputs, the Morishima elasticity is the appropriate measure of substitutability, and is not guaranteed to be symmetric unless the production function is CES (Blackorby and Russell, 1989).

Under the nested CES with emissions in the outer nest, all three Morishima elasticities should equal the outer-nest elasticity σ_n:

$$M_{KE} = M_{LE} = M_{ME} = \sigma_n$$

If instead capital, labor, and materials had different substitutability with emissions, we would observe $M_{KE} \neq M_{LE} \neq M_{ME}$.

Estimation. To test whether these predictions hold in the data, we estimate M_{jE} for each input $j \in \{K, L, M\}$ using the following regression:

$$\log \frac{j_{ft}}{E_{ft}} = M_{jE} \log P_t^E + \sum_{k \neq E} \beta_{jk} \log P_{k,t} + \gamma_f + h_n(t) + u_{ft} \tag{40}$$

where f indexes firms, n indexes sectors, and t indexes years. We include firm fixed effects γ_f to control for time-invariant firm characteristics and sector-specific cubic time trends $h_n(t)$ to capture systematic declines in emissions intensity not explained by prices or input mix—what we later interpret as clean technology B.

The key identification challenge is that $\log(j_{ft}/E_{ft})$ and $\log P_t^E$ are simultaneously determined in equilibrium. We address this using an instrumental variables strategy, instrumenting for the carbon price using high-frequency carbon price surprise series from Känzig (2025) and Jezequel-Royer and Levieuge (2025), aggregated to annual frequency. These surprises capture exogenous information shocks (e.g., court rulings, parliament votes) that shift the carbon permit supply.

Results. Table A3 reports our estimates of the three Morishima elasticities. All three are uniformly low (below 0.1) and statistically indistinguishable from each other. This pattern supports the nested CES assumption: emissions appear to be similarly complementary with

all productive inputs, consistent with treating capital, labor, and materials as a composite in the outer nest.

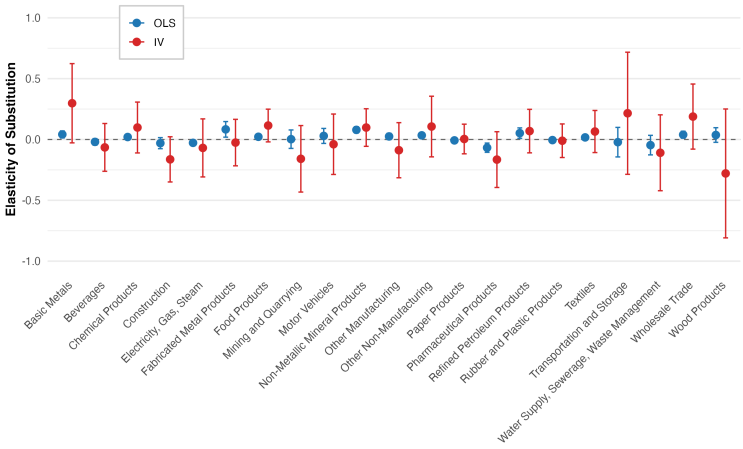
Table A3: Morishima Elasticities of Substitution

	$\log\left(\frac{K}{E}\right)$		$\log\left(\frac{L}{E}\right)$		$\log\left(\frac{M}{E}\right)$	
Model:	(1)	(2)	(3)	(4)	(5)	(6)
$\log(P_E)$	0.0146	0.0246	0.0165	0.0237	-0.0038	0.0151
	(0.0063)	(0.0544)	(0.0052)	(0.0497)	(0.0067)	(0.0525)
	OLS	IV	OLS	IV	OLS	IV
Input Price Controls	Yes	Yes	Yes	Yes	Yes	Yes
Time trend	Cubic	Cubic	Cubic	Cubic	Cubic	Cubic
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	38,786	38,786	38,786	38,786	38,786	38,786

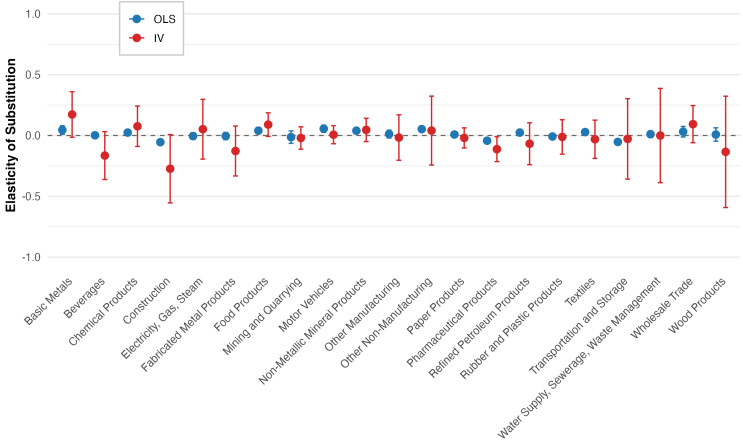
Notes: Standard errors (two-way clustered by firm and year) in parentheses. All specifications include firm fixed effects, cubic sector-specific time trends, and controls for other input prices. IV regressions instrument for $\log(P^E)$ using carbon price surprises from Känzig (2025) and Jezequel-Royer and Levieuge (2025).

Sector Heterogeneity. Figure A6 shows that the low elasticity result holds across sectors. While there is some heterogeneity in point estimates, the IV estimates are uniformly close to zero and statistically indistinguishable across all major emitting sectors. This suggests that the near-Leontief relationship between emissions and other inputs is a general feature of production in regulated industries, not driven by any particular sector.

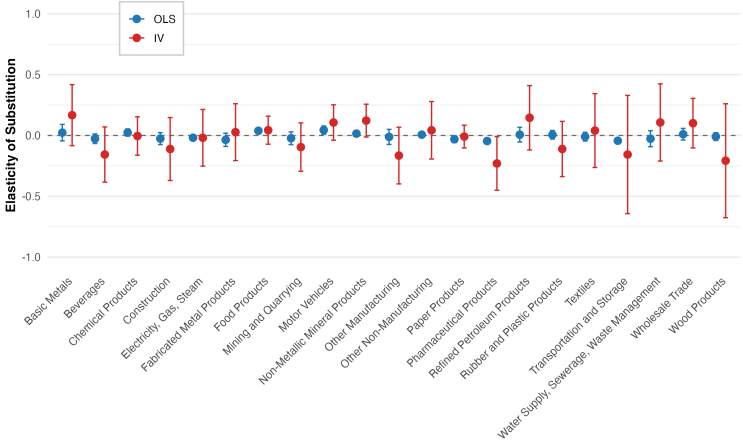
Figure A6: Morishima Elasticities by Sector



(a) Capital-Emissions (M_{KE})



(b) Labor-Emissions (M_{LE})



(c) Materials-Emissions (M_{ME})

Notes: Each panel shows sector-level estimates of the Morishima elasticity of substitution between the indicated input and emissions. Blue circles are OLS estimates; red diamonds are IV estimates using carbon price surprise instruments from Känzig (2025) and Jezequel-Royer and Levieuge (2025). Error bars show 95% confidence intervals based on two-way clustered standard errors. Dashed line at zero for reference.

Constant Elasticity Test. The nested CES also implies that the elasticity of substitution should be constant—that is, it should not vary with the level of relative prices. Table A4 tests this by including a quadratic term $\log(P^E)^2$ in our Morishima regressions. In both OLS and IV specifications, the coefficients on the quadratic term are small and statistically insignificant, consistent with a constant elasticity of substitution.

Table A4: Test for Constant Elasticity of Substitution

Panel A: OLS			
	$\log\left(\frac{K}{E}\right)$	$\log\left(\frac{L}{E}\right)$	$\log\left(\frac{M}{E}\right)$
Model:	(1)	(2)	(3)
$\log(P_E)$	0.137	0.033	0.032
	(0.092)	(0.064)	(0.068)
$\log(P_E)^2$	-0.032	-0.004	-0.009
	(0.024)	(0.017)	(0.018)
Observations	38,786	38,786	38,786
Panel B: IV			
	$\log\left(\frac{K}{E}\right)$	$\log\left(\frac{L}{E}\right)$	$\log\left(\frac{M}{E}\right)$
Model:	(1)	(2)	(3)
$\log(P_E)$	-0.103	-0.045	-0.164
	(0.401)	(0.342)	(0.410)
$\log(P_E)^2$	0.028	0.015	0.039
	(0.100)	(0.085)	(0.102)
Observations	38,786	38,786	38,786

Notes: Standard errors (two-way clustered by firm and year) in parentheses. All specifications include firm fixed effects and cubic time trends. IV regressions instrument for $\log(P_E)$ using emissions price surprises (JLR). First-stage F-statistics reported for IV specifications.

Summary. The evidence presented here supports our choice of the nested CES specification: Morishima elasticities are uniformly low, statistically indistinguishable across input pairs, and approximately constant across the range of observed prices. These findings justify treating capital, labor, and materials as a composite in the outer nest along with emissions.

We note two caveats. First, our estimates reflect short- to medium-run substitution possibilities. Over longer horizons, firms may have greater flexibility to restructure production processes or adopt new technologies that allow for different substitution patterns. Second, the nested CES remains a parametric restriction. While our tests fail to reject the specification, we cannot rule out more flexible functional forms that might fit equally well. Nevertheless, the nested CES offers a tractable and empirically supported framework for the analysis in the main text.

D Model Quantification Details

Here we describe how we quantify the remaining parameters of our model and then solve and simulate the counterfactuals. To do so, we recast the production side of model in normalized form. The inputs and output are normalized in a way to facilitate quantification of the model and the normalization choices will be described below.

D.1 Demand

Consumer Price Indices Demand is identical to the description in the main text. Let P_{int} be the price charged for the sector n good sold by firm i in year t and assume that the representative consumer takes these prices as given. Given the consumer's CES preferences for each nest, the price index for an sector n bundle can be shown to be:

$$P_{nt} = \left(\sum_{i \in I_n} D_{in} P_{int}^{1-\epsilon_n} \right)^{\frac{1}{1-\epsilon_n}} \tag{41}$$

while the price index for the aggregate bundle across all sectors can be shown to be:

$$\mathcal{P}_t = \left(\sum_{n \in N} D_n P_{nt}^{1-\upsilon} \right)^{\frac{1}{1-\upsilon}}. \tag{42}$$

D.2 Supply

We express outer and inner nests of the firm production functions in normalized form. The outer nest is given by:

$$Y_{int} = A_{int} Y_{in0} \left[(1 - \pi_{in}) X_{int}^{\frac{\sigma_n - 1}{\sigma_n}} + \pi_{in} \left(B_{int} \frac{E_{int}}{E_{in0}} \right)^{\frac{\sigma_n - 1}{\sigma_n}} \right]^{\frac{\sigma_n}{\sigma_n - 1}}. \tag{43}$$

The inner nest is:

$$X_{int} = \left[(1 - \delta_{in} - \omega_{in}) \left(\frac{K_{int}}{K_{in0}} \right)^{\frac{\xi_n - 1}{\xi_n}} + \delta_{in} \left(\frac{L_{int}}{L_{in0}} \right)^{\frac{\xi_n - 1}{\xi_n}} + \omega_{in} \left(\frac{M_{int}}{M_{in0}} \right)^{\frac{\xi_n - 1}{\xi_n}} \right]^{\frac{\xi_n}{\xi_n - 1}}. \tag{44}$$

From the firm's cost minimization problem, we can construct the marginal cost of producing the capital-labor aggregate in the inner nest and output in the outer nest. The unit cost of the capital-labor-materials aggregate X_{int} is:

$$P_{int}^X = \left[(1 - \delta_{in} - \omega_{in})^{\xi_n} \left(K_{in0} P_{int}^K \right)^{1 - \xi_n} + \delta_{in}^{\xi_n} \left(L_{in0} P_{int}^L \right)^{1 - \xi_n} + \omega_{in}^{\xi_n} \left(M_{in0} P_{int}^M \right)^{1 - \xi_n} \right]^{\frac{1}{1 - \xi_n}} \tag{45}$$

while the marginal cost of output Y_{int} is:

$$MC_{int} = \left[(1 - \pi_{in})^{\sigma} P_{int}^{X\,1 - \sigma_n} + \pi_{in}^{\sigma_n} \left(\frac{E_{in0} P_t^E}{B_{int}} \right)^{1 - \sigma_n} \right]^{\frac{1}{1 - \sigma_n}} \frac{1}{A_{int} Y_{in0}} \tag{46}$$

D.3 Equilibrium Conditions

Given CES demand across firms, in equilibrium, monopolistically competitive firms set output prices P_{int} as a constant markup over marginal costs:

$$P_{int} = \left(\frac{\epsilon_n}{\epsilon_n - 1} \right) MC_{int} \tag{47}$$

Within-nest demand given CES preferences is: $Q_{int} = \phi_{in} \left(\frac{P_{int}}{P_{nt}} \right)^{-\epsilon_n} Q_{nt}$, and multiplying and dividing by P_{nt} and P_{int} yields the expenditure version of within-nest demand:

$$S_{int} = \phi_{in} \left(\frac{P_{int}}{P_{nt}} \right)^{1 - \epsilon_n} S_{nt} \tag{48}$$

where S_{int} is firm i's revenue in sector n and year t, and S_{nt} is total revenue in sector n and year t.

Finally, output markets clear so that $Y_{int} = Q_{int}$.

D.4 Quantification of Remaining Model Parameters

We now describe how we quantify the remaining parameters in the model.

Demand Shares We cannot separately identify the demand share parameters from firm Hicks-neutral productivity so we assign equal demand shares across firms and across sectors:

$$D_{in} = 1/F_n \qquad D_n = 1/N \quad \text{for all } i = 1, \ldots, F_n \text{ and } n = 1, \ldots, N. \tag{49}$$

Production Function Shares We estimate the production function share parameters using input cost shares. The labor share parameter δ_{in} is given by:

$$\delta_{in} = \frac{P^L_{in0} L_{in0}}{P^L_{in0} L_{in0} + P^K_{in0} K_{in0} + P^M_{in0} M_{in0}}, \tag{50}$$

the materials share parameter ω_{in} is given by:

$$\omega_{in} = \frac{P^M_{in0} M_{in0}}{P^L_{in0} L_{in0} + P^K_{in0} K_{in0} + P^M_{in0} M_{in0}}, \tag{51}$$

while the emissions share parameter π_{in} is given by:

$$\pi_{in} = \frac{P^E_0 E_{in0}}{P^L_{in0} L_{in0} + P^K_{in0} K_{in0} + P^M_{in0} M_{in0} + P^E_0 E_{in0}}. \tag{52}$$

Emissions-Augmenting Technology We can back out emissions-augmenting technology by inverting the relative first-order condition between the productive composite and emissions to get:

$$B_{int} = \left[\frac{P^E_t / P^E_0}{P^X_{int} / P^X_{in0}} \right] \left[\frac{1 - \pi_{in}}{\pi_{in}} \frac{(P^E_t E_{int}) / (P^E_0 E_{in0})}{(P^X_{int} X_{int}) / (P^X_{in0} X_{in0})} \right]^{\frac{1}{\sigma_n - 1}} \tag{53}$$

Hicks-Neutral Productivity and Initial Output Quantities CES demand implies that we can back out prices from observed revenue shares:

$$P_{nt} = [N(S_{nt}/\mathcal{S}_t)]^{\frac{1}{1-v}} \tag{54}$$

$$P_{int} = P_{nt}\left[F_n(S_{int}/\mathcal{S}_{nt})\right]^{\frac{1}{1-\epsilon_n}} \tag{55}$$

Once we have prices, we next combine firm marginal cost in equation (46) with the firm's optimal pricing rule in equation (47), invert it to solve for $\tilde{A}_{int}\tilde{Y}_{in0}$:

$$\tilde{A}_{int}\tilde{Y}_{in0} = \frac{1}{\tilde{P}_{int}}\left(\frac{\epsilon_n}{\epsilon_n - 1}\right)\left[(1-\pi_{in})^{\sigma} P_{int}^{X\,1-\sigma_n} + \pi_{in}^{\sigma_n}\left(\frac{E_{in0}P_t^E}{B_{int}}\right)^{1-\sigma_n}\right]^{\frac{1}{1-\sigma_n}}. \tag{56}$$

Since $\tilde{A}_{int}\tilde{Y}_{in0}$ never appear separately we do not need to separately identify them. With all of these parameters estimated, we have fully specified the baseline of the model.

D.5 Simulation of Counterfactuals

We use the following steps to simulate the counterfactuals. Given some counterfactual $P_t^{E\prime}$:

1. Update firm marginal costs MC'_{int} using equation (46).
2. Using marginal costs and the estimated markups, update firm prices: $P'_{int} = \frac{\epsilon_n}{\epsilon_n - 1}MC'_{int}$.
3. Use equation (41) and (42) to update the industry and aggregate price indices, P'_{nt} and $\mathcal{P}'_t$.
4. Given aggregate expenditures and thus aggregate firm revenues are held fixed, compute industry revenues by inverting equation (54) and solving for S'_{nt}, and then compute firm revenues by inverting (55) and solving for S'_{int}.
    - Given revenues and prices, we can compute industry and firm quantities as $Y'_{nt} = S'_{nt}/P'_{nt}$ and $Y'_{int} = S'_{int}/P'_{int}$.
5. Dividing the counterfactual version of equation (7) by its factual version lets us recover the counterfactual emissions-composite ratio:
    - $\frac{X'_{int}}{E'_{int}} = \frac{X_{int}}{E_{int}}\left[\frac{B'_{int}}{B_{int}}\right]^{1-\sigma_n}\left[\frac{P_{int}^{E\,\prime}}{P_{int}^{E}}\right]^{\sigma_n}$
6. Dividing the counterfactual version of equation (43) by its factual version and then exploiting that the production function is CRS let's us pull out and solve for E'_{int} and X'_{int}:

- $E'_{int} = \left[\frac{\left[\frac{Y'_{int}}{Y_{int}}\right]^{\frac{\sigma_n-1}{\sigma_n}} \left[(1-\pi_{in}) X_{int}^{\frac{\sigma_n-1}{\sigma_n}} + \pi_{in} \left(\frac{B_{int} E_{int}}{E_{in0}} \right)^{\frac{\sigma_n-1}{\sigma_n}} \right]}{\left[(1-\pi_{in}) \left(\frac{X'_{int}}{E_{int}} \right)^{\frac{\sigma_n-1}{\sigma_n}} + \pi_{in} \left(\frac{B'_{int}}{E_{in0}} \right)^{\frac{\sigma_n-1}{\sigma_n}} \right]} \right]^{\frac{\sigma_n}{\sigma_n-1}}$
- $X'_{int} = \frac{X'_{int}}{E'_{int}} E'_{int}$

7. We can then solve for each of the productive inputs using the definition of total costs in the inner nest and the fact that since relative inner nest prices are not changing, inner nest expenditure shares and the inner nest unit are not changing:

   - $K'_{int} = \frac{R_{int} K_{int}}{R_{int} K_{int} + W_{int} L_{int} + P^M_{int} M_{int}} \frac{P^X_{int} X'_{int}}{R_{int}}$
   - $L'_{int} = \frac{W_{int} L_{int}}{R_{int} K_{int} + W_{int} L_{int} + P^M_{int} M_{int}} \frac{P^X_{int} X'_{int}}{W_{int}}$
   - $M'_{int} = \frac{P^M_{int} M_{int}}{R_{int} K_{int} + W_{int} L_{int} + P^M_{int} M_{int}} \frac{P^X_{int} X'_{int}}{p^M_{int}}$.

D.5.1 Simulations Shutting Down Reallocation

E Granular Instrumental Variable Construction

This appendix provides technical details on the granular instrumental variable (GIV) strategy used to identify the inner-nest elasticity of substitution ξ_n in Section 4.1. The empirical challenge is that wages are endogenous: local wages can comove with unobserved shocks that also affect firms' input choices. We address this by constructing a GIV for local wages following Gabaix and Koijen (2024). The GIV isolates wage variation driven by idiosyncratic labor demand shocks to other, large firms in the same local labor market so that these shocks can be seen to act as residual labor supply shifters.

We define labor markets at the region-year level. Let $\mathcal{I}_{rt}$ denote the set of firms operating in region r at time t, and let $N_{rt} \equiv |\mathcal{I}_{rt}|$ be the corresponding number of firms.

Let L_{irnt} denote firm employment. Define firm employment shares within a market as

$$\varphi_{irnt} \equiv \frac{L_{irnt}}{\sum_{j \in \mathcal{I}_{rt}} L_{jrnt}} \tag{57}$$

where $\sum_{i \in \mathcal{I}_{rt}} \varphi_{irnt} = 1$. We use the following reduced form labor demand representation:

$$\log L_{irnt} = \phi \log \tilde{P}^L_{rt} + \eta_{rt} + \tilde{u}_{irnt} \tag{58}$$

where η_{rt} is an aggregate labor demand shifter common to all firms in market (r, t), and $\tilde{u}_{irnt}$ is an idiosyncratic firm-level labor demand shifter. In our nested CES model, firm labor

demand depends on additional objects (e.g. output scale and other input prices). Equation (58) should therefore be interpreted as a reduced form projection after absorbing these other components of firm labor demand into η_{rt} and $\tilde{u}_{irnt}$.

We work in first differences and define u_{irnt} as the idiosyncratic shock to labor demand growth:

$$\Delta \log L_{irnt} = \phi \, \Delta \log \tilde{P}^L_{rt} + \Delta \eta_{rt} + u_{irnt}. \tag{59}$$

Labor supply is given by an upward-sloping reduced form:

$$\log \tilde{P}^L_{rt} = \alpha \, \log L_{rt} + \varepsilon_{rt}$$

where $L_{rt} \equiv \sum_{i \in \mathcal{I}_{rt}} L_{irnt}$ and where ε_{rt} is a market-level labor supply shifter. Taking first differences of labor supply yields:

$$\Delta \log \tilde{P}^L_{rt} = \alpha \, \Delta \log L_{rt} + \Delta \varepsilon_{rt}, \tag{60}$$

Log changes in total employment are a share-weighted average of firm-level log changes:[11]

$$\Delta \log L_{rt} \; = \; \sum_{i \in \mathcal{I}_{rt}} \varphi_{ir,t-1} \, \Delta \log L_{irt}. \tag{61}$$

Substituting equation (59) into equation (61) yields aggregate labor demand growth:

$$\Delta \log L_{rt} = \phi \, \Delta \log \tilde{P}^L_{rt} + \Delta \eta_{rt} + \sum_{i \in \mathcal{I}_{rt}} \varphi_{irn,t-1} \, u_{irnt}. \tag{62}$$

Next, combining equations (60) and (62) and solving for wage changes gives:

$$\Delta \log \tilde{P}^L_{rt} = \frac{\alpha}{1 - \alpha\phi} \sum_{i \in \mathcal{I}_{rt}} \varphi_{irn,t-1} \, u_{irnt} + \frac{\alpha}{1 - \alpha\phi} \Delta \eta_{rt} + \frac{1}{1 - \alpha\phi} \Delta \varepsilon_{rt}. \tag{63}$$

Equation (63) shows that market wages respond to a share-weighted sum of idiosyncratic firm labor demand shocks, as well as regional demand and supply shocks. When the firm size distribution is concentrated, shocks to large firms with substantial employment and high $\varphi_{irn,t-1}$ can generate sizable market-wide wage movements.

Define the market-level granular shifter as the deviation of the share-weighted shock from

[11] Note that this will not necessarily be true in levels.

the equal-weighted shock following Gabaix and Koijen (2024):

$$Z_{rt} \equiv \sum_{i \in \mathcal{I}_{rt}} \left(\varphi_{ir,t-1} - \frac{1}{N_{rt}} \right) u_{irt} = \sum_{i \in \mathcal{I}_{rt}} \varphi_{irn,t-1} \, u_{irnt} - \frac{1}{N_{rt}} \sum_{i \in \mathcal{I}_{rt}} u_{irnt}. \tag{64}$$

When instrumenting for wages we use a leave-one-out version to avoid mechanical correlation between the instrument and the firm's own shock:

$$Z_{rt}^{-i} \equiv \sum_{j \in \mathcal{I}_{rt} \setminus \{i\}} \left(\varphi_{jrn,t-1} - \frac{1}{N_{rt}} \right) \widehat{u}_{jrnt}. \tag{65}$$

To construct the instrument in equation (65) we need to recover the set of firm-year specific idiosyncratic labor demand shocks. We do so using the residuals from a high-dimensional fixed effects regression designed to remove common demand components and persistent firm differences from firm labor demand:

$$\Delta \log L_{irnt} = \zeta_i + \theta_{rnt} + \widehat{u}_{irnt}, \tag{66}$$

where ζ_i are firm fixed effects and θ_{rnt} are region-sector-year fixed effects. The latter absorb market-level conditions and sector-specific local shocks and wages. This leaves $\widehat{u}_{irt}$ as the idiosyncratic component of labor demand.[12] Given $\widehat{u}_{irt}$ and lagged shares $\varphi_{ir,t-1}$, we construct Z_{rt}^{-i} from (65) for each firm-year observation.

In addition to the instrument we work to improve identification by using a measure of regional wages instead of firm average wages. Firm-level average wages in Orbis may reflect both local market conditions and firm-specific factors (e.g. skill mix of the workforce). To obtain a wage measure closer to the market wage $\log \tilde{P}_{rt}^{L}$, we extract a region-year wage component from firm wage data. Let w_{irt}^{obs} denote the observed firm average wage (e.g. wage bill per worker). We estimate:

$$\log W_{irnt}^{\text{obs}} = \omega_{rt} + \kappa_{nt} + \nu_i + e_{irnt}^{w}, \tag{67}$$

where ω_{rt} is a region-year fixed effect and our measure of the market wage, κ_{nt} is a sector-year fixed effect, and ν_i is a firm fixed effect capturing time-invariant wage premia specific to a firm. We then set

$$\log P_{rt}^{L} \equiv \omega_{rt}. \tag{68}$$

[12] Using $\Delta \log L_{irt}$ rather than $\log L_{irt}$ aligns the shock construction with equation (61).

F Marginal Abatement Cost Derivation

F.1 Firm-Level MAC Slope

This subsection derives the firm-level MAC slope by: (i) start from the emissions-composite elasticity, (ii) use the emissions cost-share identity to translate it into a level elasticity for emissions demand, and (iii) invert to obtain the MAC slope. Under CES production, we can also obtain a closed-form expression for the full firm MAC schedule $MAC_{int}(Y_{int}, E_{int})$.

Consider a firm's dual cost problem with an emissions price P_t^E and non-emissions input price P_{int}^X:

$$C_{int}(Y_{int}, P_{int}^X, P_t^E) = \min_{X_{int}, E_{int}} P_{int}^X X_{int} + P_t^E E_{int} \quad \text{s.t.} \quad Y_{int} \leq A_{int} \left[(1-\pi_{in}) X_{int}^{\frac{\sigma_n - 1}{\sigma_n}} + \pi_{in} (B_{int} E_{int})^{\frac{\sigma_n - 1}{\sigma_n}} \right]^{\frac{\sigma_n}{\sigma_n - 1}}$$

By Shephard's lemma:

$$E_{int}(Y_{int}, P_{int}^X, P_t^E) = \frac{\partial C_{int}(Y_{int}, P_{int}^X, P_t^E)}{\partial P_t^E}.$$

Recall the firm emissions cost share:

$$s_{E,int} \equiv \frac{P_t^E E_{int}}{P_t^E E_{int} + P_{int}^X X_{int}} \in [0, 1].$$

Rearranging gives:

$$\log\left(\frac{E_{int}}{X_{int}}\right) = \log\left(\frac{s_{E,int}}{1 - s_{E,int}}\right) - \log\left(\frac{P_t^E}{P_{int}^X}\right). \tag{69}$$

Next, differentiate equation (69) with respect to $\log(P_t^E / P_{int}^X)$ holding Y_{int} fixed:

$$\left.\frac{d \log(E_{int}/X_{int})}{d \log(P_t^E/P_{int}^X)}\right|_{Y_{int}} = \frac{1}{s_{E,int}(1 - s_{E,int})} \left.\frac{d s_{E,int}}{d \log(P_t^E/P_{int}^X)}\right|_{Y_{int}} - 1.$$

Combining with the definition of σ_n yields:

$$\left.\frac{d s_{E,int}}{d \log(P_t^E/P_{int}^X)}\right|_{Y_{int}} = -(\sigma_n - 1)\, s_{E,int}(1 - s_{E,int}). \tag{70}$$

Again rearrange the definition of the firm emissions cost share to write emissions as:

$$E_{int} = s_{E,int} \frac{C_{int}(Y_{int}, P_{int}^X, P_t^E)}{P_t^E}.$$

Taking logs and differentiating with respect to $\log P_t^E$ holding (Y_{int}, P_{int}^X) fixed gives:

$$\left.\frac{d\log E_{int}}{d\log P_t^E}\right|_{Y_{int},P_{int}^X} = \left.\frac{d\log s_{E,int}}{d\log P_t^E}\right|_{Y_{int},P_{int}^X} + \left.\frac{d\log C_{int}}{d\log P_t^E}\right|_{Y_{int},P_{int}^X} - 1. \tag{71}$$

Next solve for the two derivatives on the right hand side. By Shephard's lemma the second term is:

$$\left.\frac{d\log C_{int}}{d\log P_t^E}\right|_{Y_{int},P_{int}^X} = \frac{P_t^E E_{int}}{C_{int}} = s_{E,int}.$$

Since we are holding P_{int}^X fixed, we have that $d\log(P_t^E/P_{int}^X) = d\log P_t^E$, so using equation (70) we have:

$$\left.\frac{d\log s_{E,int}}{d\log P_t^E}\right|_{Y_{int},P_{int}^X} = \frac{1}{s_{E,int}}\left.\frac{ds_{E,int}}{d\log P_t^E}\right|_{Y_{int},P_{int}^X} = -(\sigma_n - 1)(1 - s_{E,int}).$$

Substituting these into equation (71) yields the firm-level elasticity of emissions demand in levels:

$$\left.\frac{d\log E_{int}}{d\log P_t^E}\right|_{Y_{int},P_{int}^X} = -\sigma_n\,(1 - s_{E,int}). \tag{72}$$

By the definition of the MAC, the firm MAC schedule is the inverse of emissions demand at the fixed (Y_{int}, P_{int}^X):

$$MAC_{int}(E_{int}) = P_t^E \quad \text{such that} \quad E_{int}(Y_{int}, P_{int}^X, P_t^E) = E_{int},$$

so we can then write the elasticity of the inverse as:

$$\left.\frac{d\log MAC_{int}(E_{int})}{d\log E_{int}}\right|_{Y_{int},P_{int}^X} = \left(\left.\frac{d\log E_{int}}{d\log P_t^E}\right|_{Y_{int},P_{int}^X}\right)^{-1} = -\frac{1}{\sigma_n\,(1 - s_{E,int})}. \tag{73}$$

The firm MAC slope depends on the within-firm substitution elasticity σ_n and the firm's emissions cost share $s_{E,int}$.

F.1.1 Closed form Firm MAC Schedule

For a firm we can also derive a closed form expression for the MAC schedule. Consider the dual problem where emissions are constrained instead of priced. In this case, the firm chooses X_{int} to minimize costs $P_{int}^X X_{int}$ subject to the feasibility constraint:

$$Y_{int} \leq A_{int}\left[(1 - \pi_{in})X_{int}^{\frac{\sigma_n - 1}{\sigma_n}} + \pi_{in}(B_{int}E_{int})^{\frac{\sigma_n - 1}{\sigma_n}}\right]^{\frac{\sigma_n}{\sigma_n - 1}}.$$

Assuming the output constraint binds, solving for X_{int} gives:

$$X_{int}(Y_{int}, E_{int}) = \left(\frac{(Y_{int}/A_{int})^{\frac{\sigma_n - 1}{\sigma_n}} - \pi_{in}(B_{int}E_{int})^{\frac{\sigma_n - 1}{\sigma_n}}}{1 - \pi_{in}} \right)^{\frac{\sigma_n}{\sigma_n - 1}}.$$

Hence the constrained cost function is:

$$C_{int}(Y_{int}, E_{int}) = P_{int}^X \left(\frac{(Y_{int}/A_{int})^{\frac{\sigma_n - 1}{\sigma_n}} - \pi_{in}(B_{int}E_{int})^{\frac{\sigma_n - 1}{\sigma_n}}}{1 - \pi_{in}} \right)^{\frac{\sigma_n}{\sigma_n - 1}}, \tag{74}$$

where we must have that $(Y_{int}/A_{int})^{\frac{\sigma_n - 1}{\sigma_n}} > \pi_{in}(B_{int}E_{int})^{\frac{\sigma_n - 1}{\sigma_n}}$. This holds for emissions levels below $\bar{E}_{int}$ which is given by:

$$\bar{E}_{int} \equiv \frac{Y_{int}}{A_{int}B_{int}\,\pi_{in}^{\sigma_n/(\sigma_n - 1)}}.$$

Differentiating equation (74) with respect to E_{int} yields the closed-form firm MAC schedule on the region $E_{int} < \bar{E}_{int}$:

$$MAC_{int}(Y_{int}, E_{int}) = P_{int}^X \frac{\pi_{in}}{1 - \pi_{in}} B_{int}^{\frac{\sigma_n - 1}{\sigma_n}} E_{int}^{-\frac{1}{\sigma_n}} \left(\frac{(Y_{int}/A_{int})^{\frac{\sigma_n - 1}{\sigma_n}} - \pi_{in}(B_{int}E_{int})^{\frac{\sigma_n - 1}{\sigma_n}}}{1 - \pi_{in}} \right)^{\frac{1}{\sigma_n - 1}}. \tag{75}$$

Next, recognize that differentiating equation (74) with respect to Y_{int} gives the marginal cost of output under the constraint:

$$MC_{int}(Y_{int}, E_{int}) \equiv \frac{\partial C_{int}(Y_{int}, E_{int})}{\partial Y_{int}} = \frac{P_{int}^X}{A_{int}(1 - \pi_{in})} \left(\frac{Y_{int}}{A_{int}} \right)^{-\frac{1}{\sigma_n}} \left(\frac{(Y_{int}/A_{int})^{\frac{\sigma_n - 1}{\sigma_n}} - \pi_{in}(B_{int}E_{int})^{\frac{\sigma_n - 1}{\sigma_n}}}{1 - \pi_{in}} \right)^{\overline{\sigma_n}}$$

Substituting this into equation (75) yields:

$$MAC_{int}(Y_{int}, E_{int}) = MC_{int} \cdot \left[\pi_{in} \left(\frac{B_{int}A_{int}E_{int}}{Y_{int}} \right)^{\frac{\sigma_n - 1}{\sigma_n}} \right] \cdot \frac{Y_{int}}{E_{int}}. \tag{76}$$

Notice that the term in square brackets is the output elasticity of emissions for our CES production function:

$$\frac{\partial Y_{int}}{\partial E_{int}} \equiv \pi_{in} \left(\frac{B_{int}A_{int}E_{int}}{Y_{int}} \right)^{\frac{\sigma_n - 1}{\sigma_n}}. \tag{77}$$

By Shepherd's Lemma and duality we then have that this is equal to the emissions cost share

so that:

$$MAC_{int}(Y^*_{int}, E^*_{int}, X^*_{int}) = MC_{int} \cdot s_{E,int} \cdot \frac{Y^*_{int}}{E^*_{int}} \tag{78}$$

where stars represent that this holds only at the firm's cost-minimizing choices. The equilibrium marginal abatement cost can be expressed as the product of the marginal cost of output, the cost share of emissions, and the emissions intensity of the firm.

G Industry-Level Marginal Abatement Cost

This appendix derives the industry-level marginal abatement cost slope. The key insight is that when firms differ in emissions intensity, an industry can reduce aggregate emissions both (i) within firms by shifting away from emissions and (ii) across firms as cleaner firms grow and dirtier firms shrink. The reallocation channel makes the industry MAC flatter than firm-level schedules.

G.1 Industry Cost Function

Consider an industry n. Firm i has cost function $C_{int}(Y_{int}, E_{int})$, the minimum cost of producing output Y_{int} while emitting E_{int}. Define the industry cost function as the minimum total cost to produce aggregate output Y_n subject to aggregate emissions E_n:

$$C_{nt}(Y_{nt}, E_{nt}) = \min_{\{Y_{int}, E_{int}\}} \sum_i C_{int}(Y_{int}, E_{int}) \quad \text{s.t.} \quad \sum_i Y_{int} = Y_{nt}, \quad \sum_i E_{int} = E_{nt}. \tag{79}$$

The Lagrangian is

$$\mathcal{L} = \sum_i C_{int}(Y_{int}, E_{int}) - \lambda \left(\sum_i Y_{int} - Y_{nt} \right) - \mu \left(\sum_i E_{int} - E_{nt} \right).$$

The first-order conditions for the cost-minimizing allocation are:

$$\frac{\partial C_{int}}{\partial Y_{int}} = \lambda \quad \forall i \tag{80}$$

$$-\frac{\partial C_{int}}{\partial E_{int}} = -\mu \quad \forall i \tag{81}$$

The first condition is that marginal costs of production are equalized across firms while the

second is that the MAC is equalized across firms.[13]

By the envelope theorem, the industry's MAC at the optimal allocation is also the shadow value on the industry emissions constraint:

$$MAC_{nt} \equiv -\frac{\partial C_{nt}}{\partial E_{nt}} = -\mu, \tag{82}$$

and is also equal to the common firm-level MAC at the optimal allocation. Although the levels are equal at the optimal allocation, the slope of the industry MAC curve differs from firm-level slopes because the industry can additionally reallocate production across heterogeneous firms.

G.2 How Demand Determines the Cost of Reallocation

Equation (79) solves for the cheapest way to produce Y_{nt} with emissions E_{nt}. One way to lower E_{nt} while holding Y_{nt} fixed is to shift output from dirtier firms to cleaner firms. The cost of this reallocation depends on how substitutable firms' products are from the consumer's perspective. Our CES demand system gives the usual expression for demand for firm output:

$$Y_{int} = Y_{nt}\left(\frac{P_{int}}{P_{nt}}\right)^{-\varepsilon_n},$$

where $\varepsilon_n > 0$ is the within-industry demand elasticity, P_{int} are firm output prices and P_{nt} is the industry price index. When ε_n is high, firm-specific varieties are close substitutes and consumers readily switch toward cleaner firms, so reallocation is less costly. When ε_n is low, firm-specific varieties are more differentiated and consumers place more value variety, making reallocation more costly. Thus ε_n governs the strength of the between-firm reallocation channel, while the micro elasticity σ_n governs within-firm substitution.

G.3 Aggregate Elasticity Derivation

To characterize the slope of the industry MAC schedule, we follow Oberfield and Raval (2021) and summarize the industry as if it had an *aggregate* elasticity $\bar{\sigma}_n$ combining within-firm substitution and between-firm reallocation. Define the aggregate elasticity as the response of the industry emissions-to-productive composite ratio to the relative price:

$$\bar{\sigma}_n \equiv -\left.\frac{d\log(E_{nt}/X_{nt})}{d\log(P_t^E/P_{int}^X)}\right|_{Y_{nt}}. \tag{83}$$

[13]Because C_{int} is decreasing in emissions E_{int}, we have $\mu < 0$.

We now derive this aggregate elasticity. We define the following three shares necessary for computing the aggregate elasticity:

- Emissions cost share in the firm:

$$s_{E,int} = \frac{P^E E_{int}}{P^E E_{int} + P^X_{int} X_{int}} \in [0,1].$$

- Firm cost share in the industry:

$$s_{int} = \frac{P^E_t E_{int} + P^X_{int} X_{int}}{P^E_t E_{nt} + P^X_{int} X_{nt}} \in [0,1] \qquad \sum_i s_{int} = 1.$$

- Emissions cost share in the industry:

$$\Gamma_{nt} = \frac{P^E_t E_{nt}}{P^E_t E_{nt} + P^X_{int} X_{nt}} = \sum_i s_{E,int}\, s_{int} \in [0,1], \qquad E_{nt} = \sum_i E_{int}, \quad X_{nt} = \sum_i X_{int}.$$

First, differentiate the industry share Γ_{nt} with respect to log relative prices $\log(P^E_t/P^X_{int})$:

$$\frac{d\Gamma_{nt}}{d\log(P^E_t/P^X_{int})} = \sum_i \left[\frac{ds_{E,int}}{d\log(P^E_t/P^X_{int})}\, s_{int} + s_{E,int}\, \frac{ds_{int}}{d\log(P^E_t/P^X_{int})} \right]. \tag{84}$$

The first term is captures within-firm substitution. Under CES production this is equal to:

$$\left.\frac{ds_{E,int}}{d\log(P^E_t/P^X_{int})}\right|_{Y_{int}} = (1-\sigma_n)\, s_{E,int}(1-s_{E,int}).$$

The second term captures across-firm reallocation. Under CES demand across firms this is equal to:

$$\frac{d\log s_{int}}{d\log(P^E_t/P^X_{int})} = (1-\varepsilon_n)(s_{E,int} - \Gamma_{nt}).$$

Firms with above-average emissions intensity lose market share when demand is elastic and the relative price of emissions increases.

Next, define the heterogeneity index:

$$\chi_{nt} = \frac{\sum_i (s_{E,int} - \Gamma_{nt})^2\, s_{int}}{\Gamma_{nt}(1-\Gamma_{nt})} \in [0,1].$$

The heterogeneity index is a measure of dispersion in emissions cost shares within an industry.

Rearranging equation (84) yields:

$$\frac{d\Gamma_{nt}}{d\log(P_t^E/P_{int}^X)} = -\Gamma_{nt}(1-\Gamma_{nt})\left[(\sigma_n-1)(1-\chi_{nt})+(\varepsilon_n-1)\chi_{nt}\right].$$

Rearranging the definition of sectoral emissions cost shares yields:

$$\log(E_{nt}/X_{nt}) = \log\left(\frac{\Gamma_{nt}}{1-\Gamma_{nt}}\right) - \log(P_t^E/P_{int}^X).$$

Taking the log of the right hand side of and differentiating with respect to log relative prices yields:

$$\bar{\sigma}_n - 1 = -\frac{1}{\Gamma_{nt}(1-\Gamma_{nt})}\frac{d\Gamma_{nt}}{d\log(P_t^E/P_{int}^X)}.$$

Insert the previous expression for $\frac{d\Gamma_{nt}}{d\log(P_t^E/P_{int}^X)}$ and solve for $\bar{\sigma}_n$:

$$\bar{\sigma}_n = (1-\chi_{nt})\sigma_n + \chi_{nt}\varepsilon_n. \tag{85}$$

The aggregate elasticity is a convex combination of the within-firm elasticity σ_n and the demand elasticity ε_n, with weight on the demand elasticity given by the heterogeneity index χ_{nt}.

G.4 Industry MAC Slope

Next we show how the aggregate elasticity $\bar{\sigma}_n$ translates into the slope of the sectoral MAC schedule. Recall that the sectoral marginal abatement cost MAC_{nt} is defined as the shadow value on a sectoral emissions constraint at level E_{nt}, holding output fixed. By standard duality, this shadow value equals the emissions price that decentralizes the cost-minimizing allocation, so the sectoral MAC schedule is the inverse of sectoral emissions demand:

$$MAC_{nt}(E_{nt}) \equiv P_t^E \quad \text{such that} \quad E_{nt}(Y_{nt}, P_{nt}^X, P_t^E) = E_{nt}.$$

To obtain the slope of the sectoral MAC schedule, again rearrange the definition of the industry emissions cost share to get:

$$\log\left(\frac{E_{nt}}{X_{nt}}\right) = \log\left(\frac{\Gamma_{nt}}{1-\Gamma_{nt}}\right) - \log\left(\frac{P_t^E}{P_{nt}^X}\right).$$

Differentiate this with respect to $\log(P_t^E/P_{nt}^X)$ holding output fixed:

$$\left.\frac{d\log(E_{nt}/X_{nt})}{d\log(P_t^E/P_{nt}^X)}\right|_{Y_{nt}} = \frac{1}{\Gamma_{nt}(1-\Gamma_{nt})}\left.\frac{d\Gamma_{nt}}{d\log(P_t^E/P_{nt}^X)}\right|_{Y_{nt}} . - 1,$$

where following the derivation of the aggregate elasticity we have that:

$$\left.\frac{d\Gamma_{nt}}{d\log(P_t^E/P_{nt}^X)}\right|_{Y_{nt}} = -(\bar{\sigma}_n - 1)\,\Gamma_{nt}(1-\Gamma_{nt}). \tag{86}$$

Next we convert this into an elasticity of emissions demand in levels. Rearranging the definition of industry emissions cost shares yields:

$$E_{nt} = \Gamma_{nt}\,\frac{C_{nt}(Y_{nt}, P_{nt}^X, P_t^E)}{P_t^E},$$

Next, take logs and differentiate with respect to $\log P_t^E$ holding (Y_{nt}, P_{nt}^X) fixed:

$$\left.\frac{d\log E_{nt}}{d\log P_t^E}\right|_{Y_{nt},P_{nt}^X} = \left.\frac{d\log\Gamma_{nt}}{d\log P_t^E}\right|_{Y_{nt},P_{nt}^X} + \left.\frac{d\log C_{nt}}{d\log P_t^E}\right|_{Y_{nt},P_{nt}^X} - 1. \tag{87}$$

Shepherd's Lemma gives us the second term is the emissions cost share:

$$\left.\frac{d\log C_{nt}}{d\log P_t^E}\right|_{Y_{nt},P_{nt}^X} = \Gamma_{nt}.$$

Holding P_{nt}^X fixed implies $d\log(P_t^E/P_{nt}^X) = d\log P_t^E$. Applying equation (86) the first term is then:

$$\left.\frac{d\log\Gamma_{nt}}{d\log P_t^E}\right|_{Y_{nt},P_{nt}^X} = \frac{1}{\Gamma_{nt}}\left.\frac{d\Gamma_{nt}}{d\log P_t^E}\right|_{Y_{nt},P_{nt}^X} = -(\bar{\sigma}_n - 1)(1-\Gamma_{nt}).$$

Substituting both terms back into equation (87) gives the elasticity of sectoral emissions demand in levels:

$$\left.\frac{d\log E_{nt}}{d\log P_t^E}\right|_{Y_{nt},P_{nt}^X} = -\bar{\sigma}_n\,(1-\Gamma_{nt}). \tag{88}$$

Because the sectoral MAC schedule $MAC_{nt}(E)$ is the inverse of the sectoral emissions-demand function $E_{nt}(Y_{nt}, P_{nt}^X, P_t^E)$, the elasticity of the inverse function implies:

$$\left.\frac{d\log MAC_{nt}(E_{nt})}{d\log E_{nt}}\right|_{Y_{nt},P_{nt}^X} = \left(\left.\frac{d\log E_{nt}}{d\log P_t^E}\right|_{Y_{nt},P_{nt}^X}\right)^{-1} = -\frac{1}{\bar{\sigma}_n\,(1-\Gamma_{nt})}. \tag{89}$$

The industry slope is inversely proportional to the industry aggregate elasticity $\bar{\sigma}_n$. The greater the aggregate elasticity, the flatter the industry MAC schedule and the cheaper it is for additional emissions reductions. This relationship is mediated by the industry's overall emissions cost share such that more emissions intensive industries have steeper aggregate MAC curves.

H Economy-Wide Marginal Abatement Cost

This appendix derives the economy-wide marginal abatement cost schedule. The key insight is that when sectors differ in emissions intensity, the economy can reduce aggregate emissions both (i) within sectors as described in Section G and (ii) across sectors as cleaner sectors expand and dirtier sectors contract. The reallocation channel makes the economy-wide MAC flatter than sector-level schedules.

H.1 Economy-Wide Cost Function

Consider an economy with sectors indexed by $n \in \mathcal{N}$. Sector n has cost function $C_{nt}(Y_{nt}, E_{nt})$, the minimum cost of producing output Y_{nt} while emitting E_{nt}. (In our model, $C_{nt}(\cdot)$ is the sector cost function that already incorporates cost-minimizing reallocation across heterogeneous firms within sector n.) Define the economy-wide cost function as the minimum total cost to produce aggregate final demand Y_t subject to aggregate emissions E_t:

$$C_t(Y_t, E_t) = \min_{\{Y_{nt}, E_{nt}\}_{n\in\mathcal{N}}} \sum_{n\in\mathcal{N}} C_{nt}(Y_{nt}, E_{nt}) \quad \text{s.t.} \quad Y_t = \left(\sum_{n\in\mathcal{N}} D_n^{1/\upsilon} Y_{nt}^{\frac{\upsilon-1}{\upsilon}}\right)^{\frac{\upsilon}{\upsilon-1}}, \quad \sum_{n\in\mathcal{N}} E_{nt} = E_t, \tag{90}$$

where $\upsilon > 0$ is the elasticity of substitution across sectoral outputs and $D_n > 0$ are CES weights.

The Lagrangian is

$$\mathcal{L} = \sum_{n\in\mathcal{N}} C_{nt}(Y_{nt}, E_{nt}) - \lambda \left(\left(\sum_{n\in\mathcal{N}} D_n^{1/\upsilon} Y_{nt}^{\frac{\upsilon-1}{\upsilon}}\right)^{\frac{\upsilon}{\upsilon-1}} - Y_t\right) - \mu \left(\sum_{n\in\mathcal{N}} E_{nt} - E_t\right).$$

The first-order conditions for the cost-minimizing allocation are:

$$\frac{\partial C_{nt}}{\partial Y_{nt}} = \lambda \frac{\partial Y_t}{\partial Y_{nt}} \quad \forall n \tag{91}$$

$$-\frac{\partial C_{nt}}{\partial E_{nt}} = -\mu \quad \forall n \tag{92}$$

The first condition characterizes optimal reallocation of production across sectors given the final-demand aggregator, while the second condition states that the MAC is equalized across sectors.[14]

By the envelope theorem, the economy-wide MAC at the optimal allocation is also the shadow value on the economy-wide emissions constraint:

$$MAC_t \equiv -\frac{\partial C_t}{\partial E_t} = -\mu, \tag{93}$$

and is also equal to the common sector-level MAC at the optimal allocation. Although the levels are equal at the optimal allocation, the slope of the economy-wide MAC curve differs from sector-level slopes because the economy can additionally reallocate production across heterogeneous sectors.

H.2 How Demand Determines the Cost of Reallocation

Equation (90) solves for the cheapest way to produce Y_t with emissions E_t. One way to lower E_t while holding Y_t fixed is to shift demand from dirtier sectors to cleaner sectors. The cost of this reallocation depends on how substitutable sectoral outputs are from the consumer's perspective. Our CES final-demand system gives the usual expression for demand for sectoral output:

$$Y_{nt} = D_n \, Y_t \left(\frac{P_{nt}}{P_t} \right)^{-\upsilon},$$

where $\upsilon > 0$ is the cross-sector demand elasticity, P_{nt} are sector output prices, and P_t is the final-good price index. When υ is high, sectoral outputs are close substitutes and consumers readily shift toward cleaner sectors, so reallocation is less costly. When υ is low, sectoral outputs are more differentiated (or more complementary) and consumers place more value on consuming a balanced bundle across sectors, making reallocation more costly. Thus υ governs the strength of the between-sector reallocation channel, while the sectoral elasticities $\bar{\sigma}_n$ govern within-sector substitution.

H.3 Aggregate Elasticity Derivation

To characterize the slope of the economy-wide MAC schedule, we follow Oberfield and Raval (2021) and summarize the economy as if it had an *aggregate* elasticity $\bar{\sigma}_t$ combining within-sector substitution and between-sector reallocation. Define the aggregate elasticity as the

[14] Because C_{nt} is decreasing in emissions E_{nt}, we typically have $\mu < 0$.

response of the economy-wide emissions-to-productive composite ratio to the relative price:

$$\bar{\sigma}_t \equiv -\frac{d\log(E_t/X_t)}{d\log(P_t^E/P_t^X)}\bigg|_{Y_t}. \tag{94}$$

We now derive this aggregate elasticity. We define the following three shares necessary for computing the aggregate elasticity:

- Emissions cost share in the sector:

$$s_{E,nt} = \frac{P_t^E E_{nt}}{P_t^E E_{nt} + P_t^X X_{nt}} \in [0,1].$$

- Sector cost share in the economy:

$$s_{nt} = \frac{P_t^E E_{nt} + P_t^X X_{nt}}{P_t^E E_t + P_t^X X_t} \in [0,1], \qquad \sum_{n\in\mathcal{N}} s_{nt} = 1.$$

- Emissions cost share in the economy:

$$\Gamma_t = \frac{P_t^E E_t}{P_t^E E_t + P_t^X X_t} = \sum_{n\in\mathcal{N}} s_{E,nt}\, s_{nt} \in [0,1], \qquad E_t = \sum_{n\in\mathcal{N}} E_{nt}, \quad X_t = \sum_{n\in\mathcal{N}} X_{nt}.$$

First, differentiate the economy-wide share Γ_t with respect to log relative prices $\log(P_t^E/P_t^X)$:

$$\frac{d\Gamma_t}{d\log(P_t^E/P_t^X)} = \sum_{n\in\mathcal{N}} \left[\frac{ds_{E,nt}}{d\log(P_t^E/P_t^X)}\, s_{nt} + s_{E,nt}\, \frac{ds_{nt}}{d\log(P_t^E/P_t^X)} \right]. \tag{95}$$

The first term captures within-sector substitution. Using the sectoral aggregate elasticities $\bar{\sigma}_n$ derived in Section G, this is equal to:

$$\frac{ds_{E,nt}}{d\log(P_t^E/P_t^X)}\bigg|_{Y_{nt}} = (1-\bar{\sigma}_n)\, s_{E,nt}(1-s_{E,nt}).$$

The second term captures across-sector reallocation. Under CES demand across sectors this is equal to:

$$\frac{d\log s_{nt}}{d\log(P_t^E/P_t^X)} = (1-\upsilon)(s_{E,nt} - \Gamma_t).$$

Sectors with above-average emissions intensity lose expenditure share when demand is elastic and the relative price of emissions increases.

Next, define the heterogeneity index:

$$\chi_t = \frac{\sum_{n\in\mathcal{N}}(s_{E,nt} - \Gamma_t)^2 \, s_{nt}}{\Gamma_t(1-\Gamma_t)} \in [0,1].$$

The heterogeneity index is a measure of dispersion in emissions cost shares across sectors.

Because sectoral within-substitution elasticities $\bar{\sigma}_n$ can differ across n, we define the within-sector substitution component as the cost-share-weighted average:

$$\bar{\sigma}_t^N = \frac{\sum_{n\in\mathcal{N}} s_{E,nt}(1-s_{E,nt})\, s_{nt}\, \bar{\sigma}_n}{\sum_{n\in\mathcal{N}} s_{E,nt}(1-s_{E,nt})\, s_{nt}}.$$

Rearranging equation (95) yields:

$$\frac{d\Gamma_t}{d\log(P_t^E/P_t^X)} = -\Gamma_t(1-\Gamma_t)\left[(\bar{\sigma}_t^N - 1)(1-\chi_t) + (\upsilon - 1)\chi_t\right].$$

Rearranging the definition of aggregate emissions cost shares yields:

$$\log(E_t/X_t) = \log\left(\frac{\Gamma_t}{1-\Gamma_t}\right) - \log(P_t^E/P_t^X).$$

Taking the log of the right hand side and differentiating with respect to log relative prices yields:

$$\bar{\sigma}_t - 1 = -\frac{1}{\Gamma_t(1-\Gamma_t)}\frac{d\Gamma_t}{d\log(P_t^E/P_t^X)}.$$

Insert the previous expression for $\frac{d\Gamma_t}{d\log(P_t^E/P_t^X)}$ and solve for $\bar{\sigma}_t$:

$$\bar{\sigma}_t = (1-\chi_t)\bar{\sigma}_t^N + \chi_t \upsilon. \tag{96}$$

The aggregate elasticity is a convex combination of the within-sector cost share-weighted elasticity $\bar{\sigma}_t^N$ and the across-sector demand elasticity υ, with weight on the demand elasticity given by the heterogeneity index χ_t.

H.4 Economy-Wide MAC Slope

Next we show how the aggregate elasticity $\bar{\sigma}_t$ translates into the slope of the economy-wide MAC schedule. Recall that the economy-wide marginal abatement cost MAC_t is defined as the shadow value on an economy-wide emissions constraint at level E_t, holding final demand fixed. By standard duality, this shadow value equals the emissions price that decentralizes the cost-minimizing allocation, so the economy-wide MAC schedule is the inverse of aggregate

emissions demand:

$$MAC_t(E_t) \equiv P_t^E \quad \text{such that} \quad E_t(Y_t, P_t^X, P_t^E) = E_t.$$

To obtain the slope of the economy-wide MAC schedule, again rearrange the definition of the economy-wide emissions cost share to get:

$$\log\left(\frac{E_t}{X_t}\right) = \log\left(\frac{\Gamma_t}{1-\Gamma_t}\right) - \log\left(\frac{P_t^E}{P_t^X}\right).$$

Differentiate this with respect to $\log(P_t^E/P_t^X)$ holding total output fixed:

$$\left.\frac{d\log(E_t/X_t)}{d\log(P_t^E/P_t^X)}\right|_{Y_t} = \frac{1}{\Gamma_t(1-\Gamma_t)} \left.\frac{d\Gamma_t}{d\log(P_t^E/P_t^X)}\right|_{Y_t} - 1,$$

where following the derivation of the aggregate elasticity we have that:

$$\left.\frac{d\Gamma_t}{d\log(P_t^E/P_t^X)}\right|_{Y_t} = -(\bar{\sigma}_t - 1)\,\Gamma_t(1-\Gamma_t). \tag{97}$$

Next we convert this into an elasticity of aggregate emissions demand in levels. Rearranging the definition of aggregate emissions cost shares yields:

$$E_t = \Gamma_t \frac{C_t(Y_t, P_t^X, P_t^E)}{P_t^E}.$$

Next, take logs and differentiate with respect to $\log P_t^E$ holding (Y_t, P_t^X) fixed:

$$\left.\frac{d\log E_t}{d\log P_t^E}\right|_{Y_t,P_t^X} = \left.\frac{d\log\Gamma_t}{d\log P_t^E}\right|_{Y_t,P_t^X} + \left.\frac{d\log C_t}{d\log P_t^E}\right|_{Y_t,P_t^X} - 1. \tag{98}$$

Shepherd's Lemma gives us the second term is the aggregate emissions cost share:

$$\left.\frac{d\log C_t}{d\log P_t^E}\right|_{Y_t,P_t^X} = \Gamma_t.$$

Holding P_t^X fixed implies $d\log(P_t^E/P_t^X) = d\log P_t^E$. Applying equation (97) the first term is then:

$$\left.\frac{d\log\Gamma_t}{d\log P_t^E}\right|_{Y_t,P_t^X} = \frac{1}{\Gamma_t}\left.\frac{d\Gamma_t}{d\log P_t^E}\right|_{Y_t,P_t^X} = -(\bar{\sigma}_t - 1)(1-\Gamma_t).$$

Substituting both terms back into equation (98) gives the elasticity of aggregate emissions

demand in levels:

$$\left.\frac{d\log E_t}{d\log P_t^E}\right|_{Y_t,P_t^X} = -\bar{\sigma}_t\left(1-\Gamma_t\right).$$

Because the economy-wide MAC schedule $MAC_t(E)$ is the inverse of the aggregate emissions-demand function $E_t(Y_t, P_t^X, P_t^E)$, the elasticity of the inverse function implies:

$$\left.\frac{d\log MAC_t(E_t)}{d\log E_t}\right|_{Y_t,P_t^X} = \left(\left.\frac{d\log E_t}{d\log P_t^E}\right|_{Y_t,P_t^X}\right)^{-1} = -\frac{1}{\bar{\sigma}_t\left(1-\Gamma_t\right)}. \tag{99}$$

The economy-wide slope is a function of the economy aggregate elasticity $\bar{\sigma}_t$ as well as the economy emissions cost share.