

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

R. Andrew Butters*

Jackson Dorsey[†]

Ivan Rudik[‡]

August 25, 2026

Abstract

How much does pricing pollution cut aggregate emissions, and at what cost? We develop a microfounded equilibrium model of firm production and abatement that can be estimated from standard firm financial and emissions data, without observing abatement effort directly. The model captures three adjustment margins: short-run input substitution, reallocation of output from dirty to cleaner firms, and long-run clean technology adoption. Applying it to the first 17 years of the European Union Emissions Trading System (EU-ETS), we find input substitution is rigid: firms cannot cut emissions in the short run without cutting output. Reallocation instead drives nearly all of the short-run aggregate emissions response, concentrated in sectors with heterogeneous emissions intensities, like power and cement, or with elastic demand, like refining. Clean technology adoption also responds causally to permit price shocks, reducing firms' emissions intensity over the five years that follow. Doubling the permit price cuts aggregate emissions 23% in the short run and 31% in the longer run after induced clean technology adoption. In the short run, output falls by 4%, primarily because carbon costs pass from upstream polluting industries like electricity to downstream producers. In the long run, output falls just 2.5%, as clean technology adoption dampens these supply-chain impacts. Carbon pricing thus delivers substantial emissions reductions with moderate output effects, and much of the economic burden falls on downstream firms through higher upstream input costs.

JEL Codes: D22, E23, L11, O33, O44, Q52, Q54, Q58

Keywords: Carbon pricing, EU-ETS, marginal abatement costs, clean technology, reallocation, directed technical change

*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

[§]Tianzi Liu and Antonia Marcheva provided excellent research assistance. We thank Jonathan Colmer, Ken Gillingham, and Devesh Raval for insightful discussions, and we thank Dan Akerberg, Geoff Barrows, Greg Casey, Sylvain Chabé-Ferret, Jaedo Choi, Antoine Dechezleprêtre, Antton Haramboure, Emmanuel Murray-Leclair, Ezra Oberfield, Hélène Ollivier, Will Rafey, and seminar participants at the AEA Annual Meetings, NBER Economics of Industrial Decarbonization, NBER EEE, TSEE, Cornell, ITAM, LSE, PSE, TSE, Tulane University, UT Austin, University of Tokyo, Keio University, UChicago Junior Workshop, IIOC, and University of Nevada-Reno for helpful comments.

1 Introduction

Pricing pollution is the textbook fix for an externality. But how do these policies achieve emissions reductions, and at what cost? The answer depends on how firms and markets adjust in equilibrium when pollution is priced. Firms may adjust production processes, make longer-run investments to reduce emissions, while market forces may shift production across firms.

Quantifying these adjustments has remained difficult, even as pollution pricing has proliferated.¹ Two obstacles stand in the way. First, abatement effort is usually unobserved: standard firm data show emissions and input use, but not how inputs are split between production and abatement. Second, even once firm-level responses are estimated, they must still be aggregated to evaluate a policy’s total effect. Equilibrium forces drive reallocation across firms and propagate cost shocks through the production network, shaping aggregate and distributional outcomes. We develop a microfounded model of firm production and abatement that overcomes each of these obstacles. We show how to estimate the model from standard firm data without directly observing abatement effort, and how the model translates firm-level responses into aggregate equilibrium outcomes.

Our framework captures three adjustment margins through which pollution pricing reduces aggregate emissions. The first is firms’ short-run *input substitution*: firms may change which materials they purchase or which types of workers they hire in order to cut emissions, for example, by shifting electricity production within the firm from coal to gas boilers already in place. The second is *reallocation* of output from dirty to cleaner firms: pollution pricing increases output prices of dirty firms relative to clean firms, so consumers substitute toward less emissions-intensive producers. The third is firms’ long-run *clean technology adoption*: firms respond to higher pollution prices by investing in clean, emissions-saving technology, such as building new gas plants to retire existing coal plants. Clean technology accumulates over years due to time-to-build frictions, gradually reducing the emissions intensity of production.

At the core of our model are firms’ production and abatement choices. In each period, given their existing production technology, firms combine capital, labor, and materials to produce output, generating emissions as a byproduct that can be reduced through costly abatement. Abatement effort is unobserved in most empirical settings.² We show that under weak assumptions about the relationship between abatement and production, 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 used in recent work to a broad class of production technologies (Copeland and Taylor, 2004; Shapiro and Walker, 2018; Hollingsworth et al., 2026). In this representation, the elasticity of substitution between emissions and other inputs *is* the ease of abatement: when it is small, emissions and other inputs are near-complements, abatement is costly, and firms cannot cut emissions without cutting output.

¹Carbon pricing is now implemented in over 50 countries, covering nearly a third of global carbon emissions (World Bank, 2026).

²In our setting using common firm financial data, we observe aggregate expenditures on labor, capital, and materials, but not whether any given input is emitting or whether any given input is allocated toward abatement.

In our model, monopolistically competitive firms sell output both to consumers for final consumption and to other firms for use as inputs. This opens two equilibrium channels for aggregate abatement. First, the reallocation margin can operate directly by shifting final demand and production toward less emissions-intensive producers—both within and across sectors. Demand shifts toward cleaner producers can reduce aggregate emissions even if no individual firm’s emissions intensity changes. Second, because many regulated, highly polluting sectors are central in the production network, their higher emissions costs pass through to other downstream firms and sectors, which can induce another set of reallocation and aggregate emissions effects. Our approach to estimating these effects follows the work of Oberfield and Raval (2021).

While a firm’s production technology is fixed within a period, over time firms invest in clean technology. In our model, clean technology is a persistent emissions-saving productivity term that lets firms produce the same output with fewer emissions, or analogously, their inputs directed toward abatement become more productive. Because clean investments take years to plan, build, and bring online, this margin adjusts gradually.³


We use our model to study the first 17 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 35%, while industrial production remained relatively stable. These aggregate trends motivate the core policy questions we address: how much does carbon pricing reduce emissions, what mechanisms drive these reductions, what are the short- and long-run costs of these reductions in the aggregate and across sectors?

Figure 1 presents several patterns in the data that guide our modeling approach, and previews some of our findings. In Panel (a), the evolution of EU-ETS permit prices is plotted alongside emissions cost shares—defined as a firm’s emissions costs (the economic cost of retired emissions permits) divided by total costs. There are three key takeaways from this first panel. First, year-to-year, firm-level emissions cost shares move in lockstep with the permit price. If firms could substitute away from emissions cheaply, emissions would fall relative to other inputs as its price rose, and the emissions cost share would climb more slowly than the price. Through the lens of our model, the tight comovement implies that short-run abatement is costly: firms can still cut emissions, but only by cutting output, not by shifting their inputs toward abatement. Second, over the full 2005–2021 sample, permit prices rose by roughly 120% while average firm-level emissions cost shares rose by only approximately 50%, much less than one-for-one. The longer-run decoupling of prices and firm-level emissions cost shares suggests firms reduced their emissions intensity gradually over time. Third, the economy-wide aggregate emissions cost share only weakly tracks permit prices, indicating that the market as a whole responds to higher prices by reallocating production toward less emissions-intensive firms.⁴ Panel (b) substantiates this last point by showing how the sales

³For example, Colmer et al. (2024) find that French firms regulated under EU-ETS upgraded their capital stock, lowering emissions intensity over the several years following the introduction of the cap-and-trade program.

⁴Note that over the full sample, the aggregate emissions share of costs changes by only 0.3pp, compared to a 2.2pp

Figure 1: Permit Prices, Emissions Cost Shares, and Sales Shares by Emissions Intensity



Notes: Panel (a): The solid blue line shows the average annual permit price. The dash-dotted orange line depicts the mean firm-level emissions cost share—a firm’s emissions multiplied by the permit price, divided by its total input cost (the sum of capital, labor, materials, and emissions costs). The long-dashed red line shows the aggregate emissions cost share, calculated using total emissions costs and total input costs across all firms. Panel (b): the sales share of the firms in each decile of the 2005 emissions intensity distribution. Both panels use a balanced panel of firms.

share of the firms that fall into each decile of the 2005 emissions intensity distribution evolves over time. The cleanest decile of firms (i.e., decile 1) has an increasing sales share, while the other deciles have shrinking or stable sales shares.⁵

While our theoretical framework can accommodate a broad class of production technologies, our quantitative model must capture the three margins reflected in these trends: limited short run abatement, long run within-firm decarbonization, and reallocation from dirty to clean firms. We show that a nested constant elasticity of substitution (CES) specification fits the production data well, and we adopt this functional form for our baseline analysis. The inner nest combines capital, labor, and materials into a productive composite while the outer nest combines that composite with emissions. Clean technology enters the outer nest as an emissions-saving productivity term. The outer nest substitution elasticity governs the ease of abatement in the short run via input substitution. Clean technology determines the emissions that result from a given choice of productive inputs. In our model, firms take clean technology as given in any period, and endogenously invest to improve clean technology in the long run. We specify the demand side of our model as nested CES across firms within a sector, and across sectors. Nested CES allows for dirty firms or dirty sectors as a whole to shrink in response to carbon pricing while cleaner firms and sectors may grow.

Our estimates of the model parameters and counterfactual simulations reveal four key findings. First, the input substitution margin is rigid: firm-level elasticities of substitution between emissions and the productive composite are essentially zero for all sectors. Holding technology fixed, firms increase in the average firm-level emissions share of costs.

⁵The patterns shown in Panel (b) are robust to whether we use different years to define the emissions intensity deciles. The patterns in both panels are robust to whether we instead use the full sample of firms in our EUTL-Orbis matched dataset.

cannot reduce emissions in the short run without cutting output. We identify this elasticity from the firm’s cost-minimizing first-order conditions using permit supply shocks captured in the carbon price surprise series of Känzig (2025), and extended by Jézéquel-Royer and Leveuge (2025).

Second, reallocation across firms is a major driver of the aggregate emissions response to carbon pricing. Its strength depends on two parameters. The first is an emissions heterogeneity index, which captures within-sector dispersion in clean technology and how much a common carbon price disperses unit costs and output prices across firms.⁶ The second is a sector-specific demand elasticity, which governs how aggressively market shares shift toward cleaner firms as those relative prices change. Both vary widely across sectors: the heterogeneity index is an order of magnitude larger in power than in paper products, and demand elasticities range from about 4 in rubber and plastics to 12 in refining. The short-run elasticity of aggregate emissions to the carbon price is -0.34 , but collapses to -0.05 if reallocation is shut down, confirming that reallocation drives most of the short-run response. The largest reductions are concentrated in power, cement, and refining: the first two through substantial heterogeneity in clean technology, refining through highly elastic demand.

Third, clean technology responds endogenously to carbon prices, but with a lag. The average firm’s clean technology improved by approximately 46% over our sample period (Appendix Figure A8), but this observed trend reflects many drivers—fossil fuel prices, exogenous cost reductions in clean technology over time, and other regulatory mandates as well as carbon pricing. To isolate the causal contribution of the carbon price, we estimate the dynamic response of clean technology to permit price shocks via local projections, using the same carbon price surprise series as the source of identifying variation in carbon prices. We find that increases in the carbon price induce firms to shift investment toward cleaner technologies: pooling across firms, the five-year permit price elasticity of clean technology is approximately 0.211—meaning a 1% permit price increase raises firms’ clean technology by roughly 0.211% within five years—increasing to 0.318 when sectors are weighted by their share of emissions, because the largest clean technology responses are in carbon-intensive sectors such as refining and power generation.

Fourth, our counterfactuals show carbon pricing moves emissions several times more than it moves overall economic activity. Doubling the permit price cuts aggregate EU-ETS emissions by 23% in the short run while output falls by only 3.6%. After five years of induced clean technology adoption, the emissions cut deepens to 31%⁷ while easing the output loss to 2.5%. The asymmetry has a clear source. Reallocation of output toward cleaner firms in direct response to emissions pricing affords deep emissions reductions with moderate aggregate output declines. The output declines that do occur are almost entirely because upstream carbon costs pass down the production network to downstream firms, even clean ones.⁸ For example, a higher carbon price raises electricity

⁶In general, the strength of reallocation depends on heterogeneity in firms’ emissions cost shares, which are also a function of other production parameters. Since we find that emissions are statistically perfect complements with productive inputs, the cost share is solely a function of clean technology and factor prices.

⁷This is less than the 39% a naïve sum of the two margins would suggest, because the 16pp contributed by clean technology adoption on its own partly substitute for reallocation.

⁸We also find limited leakage to non-ETS firms. Because non-ETS emissions are unobserved, we impute them

prices which increases input costs for downstream manufacturing firms, thus reducing their output.

These effects vary sharply across sectors. Economic costs are highest where demand is inelastic (cement, power) or where exposure to upstream costs passed-through the production network is greatest (metals, chemicals). Sectors less exposed through the production network, such as paper, cut emissions with less effect on output or prices. Together, the three margins reconcile a rigid short-run response at the firm level with large aggregate emissions reductions and moderate impacts on output and employment.

Our results resolve a tension between two strands of the carbon pricing literature that have proceeded in parallel. A large applied microeconomic literature uses quasi-experimental designs to identify the firm-level effects of carbon pricing on 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). Studies of firm output generally find small or null impacts (Martin et al., 2016; Dechezleprêtre et al., 2023; Colmer et al., 2024). A more recent macroeconomic literature uses time series methods to estimate the aggregate effects of carbon pricing (Metcalf and Stock, 2023; Bjørnland et al., 2024; Känzig, 2025; Barrutiabengoa Ortubai et al., 2025). Känzig (2025), in particular, finds that carbon price shocks generate sizable output contractions. Our model shows the two are consistent because they measure different objects. The direct effect of a carbon price on a firm’s own output is modest. Aggregate output falls 6 times more in equilibrium because upstream carbon costs—from sectors like power—pass down the production network to downstream firms, raising their materials costs and further suppressing output. The micro estimates, by comparing regulated and unregulated firms, difference out much of this economy-wide component (Nakamura and Steinsson, 2018), precisely what the macro estimates capture. Thus, our framework builds on both literatures by bringing firm-level data to quantify aggregate equilibrium responses, which allows us to shed light on the causes of aggregate effects and the firm-level primitives that underpin them.

In addition to carbon pricing, our framework contributes to two other literatures. First, we build on a recent line of work that models emissions as an input to production in quantitative equilibrium models (Shapiro and Walker, 2018; Hollingsworth et al., 2026).⁹ We derive closed-form marginal abatement cost curves at the firm, sector, and economy-wide levels that account for equilibrium reallocation, extending the aggregation methodology of Raval (2019) and Oberfield and Raval (2021) from capital-labor substitution to emissions abatement. Second, by identifying the causal response of clean technology adoption to carbon pricing, we contribute to the directed technical change and induced innovation literatures (Popp, 2002; Acemoglu et al., 2012; Calel and Dechezleprêtre, 2016; Calel, 2020; Casey, 2024), providing firm-level evidence that carbon pricing itself redirects investment toward cleaner production technologies.¹⁰

and find leakage rates between -0.7–2.2% across five imputation strategies in the five-year counterfactual with clean technology adoption.

⁹Appendix C shows that this representation holds much more generally under standard regularity conditions, following the older literature going back to Pethig (1976); Siebert et al. (1980); Baumol and Oates (1988) and Cropper and Oates (1992).

¹⁰Unlike our findings across multiple countries from 2005 to 2021, Calel (2020) finds limited adoption of clean

Although we focus on EU-ETS in this paper, the framework is portable to other settings: wherever firm-level emissions and financial data are available, researchers can evaluate the aggregate and distributional effects of pollution pricing across a range of regulatory contexts.

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 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 that permit emission of one metric ton of CO₂. The allowances can be traded, which provides incentives for companies to reduce their emissions. Firms that cut emissions can sell surplus allowances and raise revenue, while those exceeding their limits must buy additional allowances and incur costs. Compliance is strictly monitored and regulated installations must report their greenhouse gas emissions annually. These annual reports are subject to independent third-party verification.

EU-ETS has undergone multiple phases and reforms. Phase I (2005–2007) was a pilot phase establishing baselines and market infrastructure. Phase II (2008–2012) coincided with the global financial crisis, and had depressed levels of emissions and a large surplus of allowances, depressing permit prices. In response, the EU implemented backloading and the Market Stability Reserve (MSR). Backloading delayed the auction of 900 million allowances, and the MSR, which began operation in 2019, adjusts auctioned allowance supply based on the total number of unused allowances in circulation to stabilize allowance prices. 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 one source of exogenous supply-driven variation in permit prices that we exploit for identification of firm production function parameters.

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 power generation and energy-intensive manufacturing, Phase III (2013) extended it to additional industrial sectors, notably aluminum, petrochemicals, and ammonia and nitric-acid production.

2.2 Data Sources

Here we describe our data sources. Appendix B provides more details on how we use these data to construct variables for estimation and simulation.

technology in UK firms using data through 2012. Given EU-ETS began in 2005 and our finding that clean technology takes years to adopt, these results are not inconsistent with one another.

Firm-Level Financials 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.

Materials Quantities and Prices Orbis reports materials spending but not physical quantities. We recover firm-level materials quantities by dividing reported firm-level materials costs by the country-sector-year intermediate input price index P_{nt}^M from EU KLEMS. These quantities enter our estimation of the elasticity of substitution among capital, labor, and materials, as well as the simulation. We also use these price indices as an instrument for the cross-sector demand elasticity.

Input-Output Structure and Final Demand We use Eurostat Supply-Use tables to measure two objects that enter the counterfactuals. First, intermediate use flows across sectors are used to calibrate the sector coefficients in a producer’s materials bundle. Second, final demand expenditures are used to calibrate the final demand shares across sectors for the demand system. The same final demand expenditures also serve as the outcome variable when we estimate the cross-sector demand elasticity, with the country-sector-year gross output price index from EU KLEMS as the price facing households.

Capital Rental Prices and Deflators We construct capital rental prices using a Jorgenson formula and data from EU KLEMS capital accounts and Penn World Table (PWT) 11.0 (Feenstra et al., 2015) country-year real returns and price indexes. The same PWT price information is used to express revenues, input expenditures, fixed assets, value added, the intermediate input price index, and emissions prices in real terms.¹¹

Emissions 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 entity 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).

¹¹Because the deflator is country-year specific, the real permit price can vary across countries even though the nominal allowance price is common to all firms. This variation is, however, negligible.

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 permit price for each year.

Carbon Price Surprise Series We obtain the high-frequency carbon price surprise series of Känzig (2025), extended by Jézéquel-Royer and Leveuge (2025). The series measures unexpected changes in EU-ETS permit prices around discrete regulatory announcements—European Commission auction decisions, Parliament votes, and court rulings—by computing the change in carbon futures prices in a narrow window around each event. We aggregate the daily series to annual frequency and use it as an instrument for the permit price in our estimation.

2.3 Sample Construction

We match 57% of verified EU-ETS emissions to Orbis firms, then apply common accounting data screens—chiefly requiring positive reported materials costs, which several countries’ Orbis filings omit entirely. In our primary matched EUTL-Orbis panel we have 4,671 regulated firms across 24 countries from 2005 to 2021, for which we observe both firm-level financials and verified emissions, still covering 29% of emissions under EU-ETS. This matched sample underlies our main estimation and counterfactuals. The one exception is the capital-labor-materials elasticity of substitution: because it is identified from wage variation and does not require verified emissions, we estimate it on a broader Orbis panel spanning 1998–2021 via a granular instrumental variable (GIV) design. These regressions contain 25.8 million observations in total; Appendix F describes the construction of the instrument.

We also construct a larger Orbis panel that retains all Orbis firms in our sectors regardless of whether they are linked to an ETS installation, starting from the same countries and years as our primary matched dataset. This panel contains 24 million firm-year observations. We use it in an extended version of the model to quantify how EU-ETS leads to emissions leakage to unregulated firms through reallocation of production. Because verified emissions are observed only for firms linked to EU-ETS installations, the leakage extension requires us to impute emissions for non-ETS firms. Appendix D.6 provides the full imputation procedure.

Power dominates the matched sample in terms of emissions. Together with cement and glass, refining, basic metals, and chemical products, these five largest emitting sectors account for approximately 86% of matched-sample emissions, with the remainder spread across a range of manufacturing and non-manufacturing activities. Appendix Table A1 reports the full sectoral breakdown.

Because the EU-ETS regulates *installations* based on their physical activity—combustion above 20 MW, cement production, metal ore sintering, and so on—rather than firms based on their industry classification, a firm’s Orbis NACE code may not reflect the activity of its regulated installation. For example, a construction company may operate an asphalt mixing plant, a holding company

may run a power plant, and an oil extraction firm could own a refinery. We classify firms by their primary Orbis NACE code. This captures the firm’s primary economic activity, which better represents the economic decision-making (and financial data) for the firm than the physical activity of an individual regulated installation. This yields 11 sectors: six correspond to directly regulated EU-ETS activities (power, cement and glass, basic metals, refining, chemical products, and pulp and paper), two retain distinct manufacturing identities (rubber and plastics, fabricated metals), mining and quarrying forms its own group, and the remaining sectors are consolidated into light manufacturing and other non-manufacturing.

To get a sense of whether this classification is economically appropriate, we compare firms’ Orbis sectors to the sector of the matched installations in the EUTL. For the five largest emitting sectors, a firm’s 2-digit Orbis-reported NACE sector and the 2-digit NACE sector implied by its EUTL installations agree 87.2% of the time in chemicals at the low end, up to 99.7% of the time at the high end for power when weighting by emissions. Agreement is slightly lower in the five sectors that do not correspond to a directly regulated activity (light manufacturing, mining and quarrying, rubber and plastics, fabricated metals, and other non-manufacturing). For these sectors, the Orbis sector is the economically relevant one for about three-quarters of firms. These firms operate in a genuine production sector that matches their Orbis sector, while their regulated installation is typically an on-site combustion or power unit. Classifying these firms by their installation sector would misclassify manufacturing or mining firms as power producers. The remaining cases are concentrated in our residual other-non-manufacturing category. These are typically firms whose Orbis code reflects a legal parent, holding company, or trading entity.

Appendix Figure A1 shows how sectoral emissions shares have evolved over time. While power generation has consistently dominated emissions, its share has declined from 2005 to 2021, reflecting improvements in power sector efficiency and the relative growth of other regulated sectors.

3 Model

In our model i indexes firms, n indexes sectors, and t indexes years. We use $\mathcal{N}$ to denote the set of sectors. Our model has a short run and a long run. In the short run, monopolistically competitive firms take clean technology as given and choose inputs and output prices to maximize profits. The short run captures firms’ input substitution margin and reallocation of output across firms with heterogeneous emissions intensities. In the long run, firms invest in clean technology, which accumulates through a reduced-form law of motion and alters the production technology that enters future short-run equilibria.

3.1 Production and Abatement

Firms produce output using capital, labor, and materials. Production generates emissions as a byproduct that can be reduced through costly abatement. The challenge for estimation is that abatement effort is typically unobserved. Researchers often observe firms’ emissions and their use

of capital, labor, and materials, but not how they allocate resources toward abatement activities. Our approach addresses unobserved abatement by showing that, under a weak set of assumptions on production, the underlying model is equivalent to a production function where emissions enter as an input. This equivalence means we can apply standard production function estimation techniques directly on observed inputs and emissions without needing to observe abatement. In our setting we impose a nested constant elasticity of substitution functional form.

Short-run production proceeds in two stages. First, firms combine capital, labor, and materials into a productive composite. Second, they choose how much of this productive capacity to allocate toward output versus abatement. In the first stage, firm i in sector n at time t produces the composite:

$$X_{int} = \left[\delta_{K,in} (K_{int})^{\frac{\xi_n-1}{\xi_n}} + \delta_{L,in} (L_{int})^{\frac{\xi_n-1}{\xi_n}} + \delta_{M,in} (M_{int})^{\frac{\xi_n-1}{\xi_n}} \right]^{\frac{\xi_n}{\xi_n-1}}. \quad (1)$$

This composite represents a firm's productive capacity, the total resources it can devote to output and abatement. The δ 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 productive capacity toward abatement activities rather than output. Let $\vartheta_{int} \in [0, 1)$ denote the fraction of the productive composite allocated to abatement rather than output. For example, a steel plant may substitute scrap metal for virgin ore to avoid the ore reduction that generates most of its direct emissions, but at the cost of additional capital, labor, and materials inputs for processing. Conditional on the input and abatement choice, net output is:

$$Y_{int} = (1 - \vartheta_{int}) A_{int} X_{int} \quad (2)$$

where A_{int} is Hicks-neutral productivity. This formulation is directly analogous to the “share of output devoted to abatement” approach in the literature (e.g. Copeland and Taylor, 2004; Shapiro and Walker, 2018; Hollingsworth et al., 2026).

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

$$\frac{E_{int}}{X_{int}} = \psi(\vartheta_{int}; \pi_{in}, \sigma_n, B_{int}).$$

The function ψ is decreasing and convex in ϑ_{int} and B_{int} . We adopt a functional form governed by three parameters:

$$\psi(\vartheta_{int}) = \frac{1}{B_{int}} \left(\frac{(1 - \vartheta_{int})^{\frac{\sigma_n-1}{\sigma_n}} - (1 - \pi_{in})}{\pi_{in}} \right)^{\frac{\sigma_n}{\sigma_n-1}}. \quad (3)$$

Here B_{int} is the firm's *clean technology*, σ_n governs the curvature of abatement costs, and π_{in} is a share parameter.

Combining equations (2)–(3) yields an equivalent representation of production where emissions

appear as an input to producing net output rather than a byproduct:

$$Y_{int} = A_{int} \left[(1 - \pi_{in}) X_{int}^{\frac{\sigma_n - 1}{\sigma_n}} + \pi_{in} (B_{int} E_{int})^{\frac{\sigma_n - 1}{\sigma_n}} \right]^{\frac{\sigma_n}{\sigma_n - 1}}. \quad (4)$$

In this representation, the abatement function 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 towards abatement efforts. When σ_n is low, emissions and productive inputs are near-complements and firms can only make deep emissions cuts by substantially reducing output.¹² Appendix Figure A2 illustrates this relative price logic, with Panel A2a showing the initial input choice and Panel A2b showing the response to an emissions price increase.

The share parameter π_{in} determines the baseline emissions intensity of production. In the Cobb-Douglas special case, π_{in} determines the cost share of emissions. B_{int} is an emissions-saving technology parameter that we call *clean technology*. It multiplies emissions in equation (4), so higher B_{int} means a firm can produce the same level of output with fewer emissions. For example, a firm with gas-fired power plants has higher B_{int} than a firm with coal plants but identical balance sheet information because it emits roughly half the CO₂ per kWh given the same observed levels of capital, labor, and materials.¹³

ϑ_{int} is a firm’s unobserved within-period choice of abatement intensity conditional on its existing state of technology, while B_{int} reflects predetermined differences in technology. In the long run, firms can invest in clean technology and raise B_{int} , effectively relaxing the trade-off between emissions and the productive composite use in yielding net output (Section 3.4). This division gives the model a “putty-clay” like structure (Johansen, 1959; Atkeson and Kehoe, 1999): σ_n is the *ex post* elasticity, governing the substitution/abatement possibilities available to a firm with equipment it has in place, while investment in B_{int} is the *ex ante* margin that selects among all the technology and abatement possibilities available to the firm.¹⁴

The reformulation of emissions as an input is the key to estimation: equation (4) contains quantities we observe or can recover—the productive composite and emissions—even though the underlying model involves unobserved abatement effort ϑ_{int} . In Section 4, we show how to use the firm’s conditional factor demands to estimate the nested CES production model using only factor prices and input quantities.

Although we adopt a nested CES specification, the emissions-as-input equivalence holds more generally: Appendix C proves it under standard regularity conditions—relaxing constant elas-

¹²For example, cement production releases CO₂ 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 have higher σ_n at plants that have the ability to combust and substitute amongst multiple types of fuels such as coal and gas or gas and fuel oil.

¹³Similar to a Solow residual, clean technology captures how firms with the same input expenditures on the balance sheet have different emissions intensities because of unobserved heterogeneity in the composition and emissions of these inputs.

¹⁴One important distinction that makes the connection to a “putty-clay” model not exact is that B_{int} reflects a productivity term as opposed to a capital vintage as is the case in most of these types of models.

ticities, constant returns to scale, and the separability of abatement and production—and Appendix A.1 tests the nested CES restriction with functional-form-free Morishima elasticities. We adopt the nested CES for estimation and counterfactuals because it fits the data well.

3.2 Short Run Equilibrium

We begin with the short run equilibrium of our model. Our key timing assumption is that firms' productivities A_{int} and B_{int} are predetermined and fixed at the time input choices are made. Firms can endogenously change clean technology B_{int} in response to carbon prices, but these adjustments take time and occur across periods rather than within one. We model this long-run response in Section 3.4. Within each period, firms minimize costs over inputs taking factor prices as given, set prices as constant markups over marginal cost given monopolistic competition and CES demand, and goods markets clear.

3.2.1 Firm Cost Minimization

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. Taking factor prices and technologies (A_{int}, B_{int}) as given, it chooses inputs to minimize variable costs conditional on output Y_{int} .¹⁵

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

where P_{int}^X is the CES unit cost of the composite input X_{int} . The first-order conditions imply a log-linear relationship between input ratios and relative prices. Expressing this in terms of the emissions-to-composite ratio—the natural measure of emissions intensity in our model—yields:

$$\underbrace{\log \frac{E_{int}}{X_{int}}}_{\text{Emissions Intensity}} = \underbrace{\sigma_n \log \left(\frac{\pi_{in}}{1 - \pi_{in}} \right)}_{\text{Constant}} - \underbrace{\sigma_n \log \frac{P_t^E}{P_{int}^X}}_{\substack{\text{Emissions} \\ \text{Composite} \\ \text{Price Ratio}}} + (\sigma_n - 1) \underbrace{\log B_{int}}_{\substack{\text{Clean} \\ \text{Tech}}}. \quad (5)$$

The coefficient $-\sigma_n$ is the strength of the short-run input substitution margin: it governs how firms' emissions intensity responds to a change in the relative price of emissions, holding clean technology B_{int} fixed. Higher emissions prices lower E/X , but when σ_n is small the response is limited. Clean technology also lowers E/X when emissions and the productive composite are complements ($\sigma_n < 1$).

The analogous inner nest conditions govern substitution among capital, labor, and materials,

¹⁵We omit the analogous inner nest problem of minimizing unit cost for X_{int} .

one for each input pair. Here we display the materials-to-labor ratio:

$$\underbrace{\log \frac{M_{int}}{L_{int}}}_{\text{Input Ratio}} = \underbrace{\xi_n \log \left(\frac{\delta_{M,in}}{\delta_{L,in}} \right)}_{\text{Constant}} - \underbrace{\xi_n \log \frac{P_{int}^M}{P_{int}^L}}_{\substack{\text{Materials} \\ \text{Labor} \\ \text{Price Ratio}}} . \quad (6)$$

Higher labor prices induce firms to substitute toward materials, raising M/L ; the magnitude of this response identifies ξ_n .

3.2.2 Demand and Markups

Product demand has two layers. Within each sector, a competitive sectoral composite producer bundles firm varieties into a sectoral composite using CES technology with within-sector demand elasticity ϵ_n . Households then allocate exogenously fixed final expenditure across sectoral composites through a second CES nest with cross-sector demand elasticity v and a corresponding final demand price index P_t^F . We specify the full demand system in Appendix D. Because firms also purchase sectoral composites as materials inputs (Section 3.2.3), total sector demand includes both household final demand and materials demand from downstream firms.

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

$$\mu_n = \frac{P_{int}}{MC_{int}} = \frac{\epsilon_n}{\epsilon_n - 1} . \quad (7)$$

The firm's problem can thus be decomposed into minimizing costs conditional on output, and setting price as a constant markup over marginal cost and letting the demand system determine its output level.

3.2.3 Input-Output Linkages

Input-output linkages are central in this setting because ETS-regulated sectors produce substantial amounts of materials for use by firms, and they are also users of carbon-intensive materials from other ETS-regulated sectors. For example, electricity and refined petroleum products from carbon-intensive power generators and refineries are inputs into manufacturing and cement. A carbon price that raises the costs of upstream energy sectors is passed-through to downstream manufacturers' marginal costs through their materials input costs. This amplifies and propagates the negative economic effects of the carbon price. It also induces reallocation, both within and across sectors, beyond the direct effects of carbon pricing.

We model firms' materials input in equation (1) as a Cobb-Douglas aggregate of sectoral composites:

$$M_{int} = \varrho_{in}^{-1} \prod_{m \in \mathcal{N}} \left(\frac{M_{imnt}}{\Omega_{mnt}} \right)^{\Omega_{mnt}} . \quad (8)$$

The aggregator is sector-specific with materials expenditure shares Ω_{mnt} for user sector n buying materials from supplier sector m , where $\sum_{m \in \mathcal{N}} \Omega_{mnt} = 1$. The aggregator has a firm-specific shifter $\varrho_{in} > 0$ allowing baseline materials prices to differ across firms within a sector. The implied firm-level materials unit costs are:

$$P_{int}^M = \varrho_{in} P_{nt}^M \quad \text{where} \quad P_{nt}^M = \prod_{m \in \mathcal{N}} (P_{mt})^{\Omega_{mnt}}. \quad (9)$$

3.2.4 Goods Market Clearing

Let $S_{nt} \equiv P_{nt} Y_{nt}$ denote total sector sales and let $S_{nt}^F \equiv P_{nt} Q_{nt}^F$ denote household final expenditure on sector n . Goods market clearing requires that gross sector sales equal final demand plus materials demand from downstream sectors:

$$S_{nt} = S_{nt}^F + \sum_{m \in \mathcal{N}} \Omega_{nmt} \alpha_{mt}^M S_{mt}, \quad (10)$$

where $\alpha_{nt}^M = \mu_n^{-1} \sum_i \frac{P_{int} Y_{int}}{S_{nt}} s_{M,int}$ is materials expenditure as a share of sector n sales and $s_{M,int}$ is the share of materials in firm i 's variable costs (see Appendix D).

In our main model we do not model the labor or capital markets and instead assume there is a perfectly elastic supply of both inputs. Appendix E introduces a full general equilibrium version of the model with markets for labor, capital, and emissions, and a fixed aggregate amount of each.

3.2.5 Equilibrium

An equilibrium in each period t consists of firm-level prices, quantities, and input choices $\{P_{int}, Y_{int}, X_{int}, E_{int}, K_{int}, L_{int}, M_{int}, \{M_{imnt}\}_{m \in \mathcal{N}}\}_{i,n}$, sectoral objects $\{P_{nt}, Y_{nt}, S_{nt}, P_{nt}^M\}_n$, household final demand quantities $\{Q_{nt}^F\}_n$, and the final demand price index P_t^F such that: (i) each firm minimizes costs given factor prices and technologies (A_{int}, B_{int}) ; (ii) firms set prices as constant markups over marginal cost; (iii) sectoral composite producers and households satisfy the CES demand systems in Section 3.2.2, given exogenous aggregate household expenditure; (iv) materials prices satisfy equation (9); and (v) goods markets clear according to equation (10).

Capital and labor prices are fixed and firms face a common emissions permit price P_t^E ; materials prices and gross sectoral sales are endogenous via the input-output fixed point. This equilibrium representation preserves supply-chain pass-through while abstracting from movements in primary factor prices.

3.3 Aggregate Elasticities and Marginal Abatement Costs

Equation (5) showed that σ_n governs how a single firm's emissions intensity responds to the price of emissions. Aggregate emissions also fall in the short run for a second reason: carbon pricing raises unit costs more at firms that are emissions-intensive, or use materials from dirty suppliers, compared to firms that emit little or source clean inputs. Through the demand system of Section 3.2.2, this

change in relative unit costs across firms moves output and market share toward cleaner producers. Quantifying what a carbon price achieves therefore requires an aggregate sector-level analogue to σ_n that captures both.

This subsection shows how to construct the aggregate sector-level elasticity of substitution between emissions and the productive composite $\bar{\sigma}_{nt}$ following Oberfield and Raval (2021). Unlike σ_n , it is not a technology parameter but a statistic of the sector’s composition, which is why it carries a time subscript. We then show how the elasticity determines aggregate marginal abatement cost curves.

A sector’s ease of abatement has two components: substitution away from emissions within the firm, and reallocation of output across firms. When permit-cost pass-through—how strongly the permit price raises a firm’s composite input price through its supply chain—is common across firms or uncorrelated with their emissions cost shares, Appendix H.2 shows that the aggregate elasticity is a simple convex combination of the two components:¹⁶

$$\bar{\sigma}_{nt} = \underbrace{(1 - \chi_{nt})\sigma_n}_{\text{within-firm}} + \underbrace{\chi_{nt}\epsilon_n}_{\text{reallocation}}. \quad (11)$$

A high $\bar{\sigma}_{nt}$ means the sector cuts emissions readily when the carbon price rises while a low one means the sector must shrink for emissions to fall.

The reallocation component is the product of the within-sector demand elasticity ϵ_n and the emissions heterogeneity index $\chi_{nt} \in [0, 1]$, which measures the dispersion in emissions intensity across firms. We make χ_{nt} precise below. Both factors are necessary for reallocation. Dispersion is what makes a common carbon price a *relative* price change, raising unit costs by more at dirty firms than at clean ones.¹⁷ Elastic demand is what turns that cost gap into a shift in market shares. If firms are alike, a carbon price leaves relative prices unchanged, so no shares move, regardless of how elastic demand may be. If demand is inelastic, relative prices move but buyers barely respond so production does not shift. Only when we have dispersion in emissions intensity and elastic demand does output move toward cleaner producers. In this case, sector emissions fall even without any firm changing its technology or its input mix.

The within-firm component is governed by σ_n , carrying the complementary weight $1 - \chi_{nt}$. If every firm in the sector is equally dirty, $\chi_{nt} = 0$ and $\bar{\sigma}_{nt} = \sigma_n$: every ton abated has to come from within firms. As firms grow more heterogeneous, χ_{nt} rises and opportunities for reallocation grow. σ_n sets what a firm can do on its own while χ_{nt} determines how much that ability matters for the sector versus reallocation, which is governed by ϵ_n . If firms differ and demand is sufficiently elastic ($\epsilon_n > \sigma_n$), reallocation can make the sector more elastic than any individual firm: $\bar{\sigma}_{nt} > \sigma_n$.

We now formally define χ_{nt} . Let $s_{E,int} \equiv P_t^E E_{int} / (P_t^E E_{int} + P_{int}^X X_{int})$ be firm i ’s emissions cost

¹⁶When pass-through instead covaries with firms’ emissions cost shares, the elasticity acquires a covariance correction term. Our results use this corrected elasticity (Appendix H.2). The correction is modest: across sectors, the median absolute change in the aggregate elasticity is 0.01 (9%), and the largest is 0.09 (19%) in refining.

¹⁷From Shephard’s lemma, cost shares—or emissions intensities with common permit prices in our case—tell us how costs respond to permit prices. These cost increases then get translated into consumer-facing output prices with a constant markup.

share. χ_{nt} is the cost-weighted cross-sectional variance of these shares, normalized to lie in $[0, 1]$:

$$\chi_{nt} \equiv \frac{\sum_i w_{i|n,t} (s_{E,int} - \Gamma_{nt})^2}{\Gamma_{nt} (1 - \Gamma_{nt})} \in [0, 1], \quad (12)$$

where $\Gamma_{nt} = \sum_i w_{i|n,t} s_{E,int}$ is the sector emissions cost share and $w_{i|n,t}$ is firm i 's share of total sector variable cost. The numerator is the dispersion in how emissions-intensive firms are, and the denominator is the most dispersion a sector with average intensity Γ_{nt} could display, so the index is zero when every firm has the same emissions cost share and approaches one when those shares are maximally spread.

The index also has a direct reading in terms of clean technology. Under Leontief production ($\sigma_n \rightarrow 0$)—which we find empirically (Section 5.1.2)—a first-order approximation reduces it to:¹⁸

$$\chi_{nt} \approx \underbrace{\Gamma_{nt} (1 - \Gamma_{nt})}_{\text{Sector-level cost share normalization}} \underbrace{\text{Var}_i^s(\log B_{int})}_{\text{Cross-firm dispersion in clean technology}},$$

where Var_i^s denotes the variance across firms, taken with the same cost weights as in equation (12). To a first-order, clean technology dispersion is a primary factor in driving reallocation.

The same construction applies one level up for reallocation *across* sectors. The economy-wide aggregate elasticity combines abatement within sectors with reallocation of spending across them, weighted by a cross-sector heterogeneity index $\chi_t \in [0, 1]$ that measures dispersion in emissions intensity across sectors:

$$\bar{\sigma}_t = \underbrace{(1 - \chi_t) \bar{\sigma}_t^N}_{\text{within-sector}} + \underbrace{\chi_t v}_{\text{cross-sector reallocation}}, \quad (13)$$

where $\bar{\sigma}_t^N$ aggregates the sectoral substitution elasticities through the input-output network and v is the cross-sector demand elasticity. Emissions and costs here are embodied: each supply chain's emissions are attributed to the final good containing them, so a sector's carbon exposure includes what its suppliers pass down.¹⁹

From elasticities to marginal abatement costs These elasticities also describe the marginal abatement cost curves commonly used in environmental policy analysis. The firm, sector, or economy-wide marginal abatement cost (MAC) is the shadow cost of cutting emissions while holding output and clean technology fixed at the corresponding level of aggregation. At the economy-wide level the fixed aggregate is household final demand (Appendix H.3). We show in Appendix H that the elasticity of a MAC with respect to emissions is inversely proportional to the elasticity of

¹⁸Appendix H.2.5 derives the approximation. The approximation error is governed by the dispersion of $\log B_{int}$, not by σ_n , and by omitted cross-firm variation in composite prices P_{int}^X , production function weights π_{in} , and base-year emissions E_{in0} .

¹⁹Appendix H.3 derives the exact economy-wide response used in our results: the within-good weight gains a covariance between embodied emissions shares and pass-through, and a markup covariance enters the response. At the observed allocation both corrections are negligible, moving the response by 0.1% relative to equation (13).

substitution. In cases where σ_n is small, firm MAC schedules are close to vertical. On the other hand, higher reallocation potential (via elastic demand and heterogeneous firms) increases $\bar{\sigma}_{nt}$ and flattens the sector MAC. Steep firm MAC schedules can co-exist with flat aggregate ones.²⁰

MAC curves are often built from the bottom up, by stacking plant- or technology-level abatement costs, or by treating a sector as a representative firm. This omits reallocation and the cost propagation that arises through the input-output network. Our MAC construction generalizes these standard approaches by incorporating reallocation and supply chain impacts. The underlying elasticities are also estimated from firm-level data rather than assumed from technology costs.

3.4 Long Run Clean Technology Dynamics

Section 3.2 explained how firms choose inputs given predetermined clean technology B_{int} . Over longer horizons, firms invest to change B_{int} , altering future short run equilibria. Rather than fully specifying the firm's dynamic investment problem, we adopt a flexible semiparametric approach that traces out how clean technology evolves in response to an exogenous change in the carbon price. A long run equilibrium is the short-run model of Section 3.2 with B_{int} replaced by this estimated path.²¹

Under a time-to-build assumption, investment $I_{in,t-1}^B$ chosen at date $t-1$ raises clean technology in period t :

$$\log B_{int} = \log B_{in,t-1} + I_{in,t-1}^B, \quad (14)$$

so the cumulative change in $\log B$ from $t-1$ through $t+h$ is the sum of interim investments:

$$\log B_{in,t+h} - \log B_{in,t-1} = \sum_{r=0}^h I_{in,t-1+r}^B. \quad (15)$$

We model firm investment with a linear reduced-form rule in which the date $t-1$ carbon price and composite input price shift clean technology investment around a firm-specific steady state rate $\bar{I}_{in}^B$:

$$I_{in,t-1}^B = \bar{I}_{in}^B + \beta_n^E \log P_{t-1}^E + \beta_n^X \log P_{in,t-1}^X + \eta_{in,t-1}, \quad (16)$$

where $\eta_{in,t-1}$ is a mean-zero residual and β_n^E, β_n^X are sector-specific reduced form coefficients. This approach is an alternative to specifying the firm's full dynamic investment problem. Substituting equation (16) into equation (15) and projecting on the date $t-1$ prices yields the reduced form local projection for clean technology, where the two-regressor form treats expectations of future $\log$

²⁰Generally, this is true only if demand at each level is more elastic than the substitution margin it aggregates, which aligns with our estimates.

²¹Appendix I derives the linear rule as a first-order approximation to the policy function of a general dynamic investment problem, under linear price forecasts, and shows that the theory of directed technical change signs its coefficient given our estimated σ_n . A one-period implementation lag and no depreciation of the clean technology stock deliver both the log-additive law of motion and the local projection specification below.

prices as linear in the date- $t - 1$ log prices and the investment residuals as orthogonal to them:

$$\mathbb{E}_{t-1}[\log B_{in,t+h} - \log B_{in,t-1}] = \alpha_{in}^{(h)} + \Psi_{h,n}^E \log P_{t-1}^E + \Psi_{h,n}^X \log P_{in,t-1}^X. \quad (17)$$

The horizon- h coefficient $\Psi_{h,n}^E$ is the cumulative impulse response of clean technology to a permit price shock, combining the direct effect of $\log P_{t-1}^E$ on $I_{in,t-1}^B$ with the indirect effect through expectations over future prices and subsequent investment. Section 4.3 describes how we estimate the reduced-form coefficients $\{\Psi_{h,n}^E\}$ using an observable empirical proxy for B_{int} .²²

The estimated path of clean technology also shifts the MAC schedules of Section 3.3. Because a higher B_{int} lets a firm produce the same output with fewer emissions, it lowers that firm's emissions cost share $s_{E,int}$ at any given permit price; we derive the closed-form firm MAC in Appendix H.1 and show there that cleaner technology shifts the fixed-output schedule downward. Aggregating across firms and sectors, clean technology investment therefore shifts the level of the economy-wide MAC schedule downward while leaving its elasticity, which equation (13) governs, approximately unchanged, as Panel (b) of Figure 5 in Section 5.4 shows.

3.5 Extension: Leakage to Unregulated Firms

Our main model focuses on EU-ETS-regulated firms since emissions at non-ETS firms are unobserved. Since unregulated firms compete with regulated ones in the same goods markets, carbon pricing may shift production toward unregulated producers and undo some of the intended emissions reductions. This phenomenon is called carbon leakage. We quantify an extended version of the model using the full Orbis set of firms to measure the extent of carbon leakage. Unregulated firms face the same demand system and input-output network as those under EU-ETS, but do not pay the permit price on their emissions, so their privately optimal abatement choice is zero and output reduces to $Y_{int} = A_{int}X_{int}$. They are still affected by policy indirectly: materials costs rise through upstream pass-through, while demand shifts toward them as regulated competitors become more expensive. These indirect effects change production at non-ETS firms in our counterfactual simulations, which we interpret as leakage. Since emissions at these firms are unobserved, we impute emissions intensities using several different strategies.

4 Estimation

We organize the estimation around the three margins developed in Section 3. Section 4.1 estimates the production elasticities ξ_n and σ_n that identify the short-run input substitution margin. Section 4.2 recovers the emissions heterogeneity index χ_{nt} and the demand elasticities ϵ_n and ν . χ_{nt} and ϵ_n pin down the reallocation margin and the marginal abatement cost curves at the sector-

²²Appendix K derives the long-run elasticity of substitution between emissions and the productive composite. When emissions and the composite are Leontief like we estimate below, the long-run elasticity of substitution is just the elasticity of clean technology with respect to the permit price $\Psi_{h,n}^E$ adjusted for the extent of pass-through of permit prices through the input-output network.

level, while v comes in at the economy-wide level. Section 4.3 uses local projections to estimate the cumulative clean technology impulse response $\Psi_{h,n}^E$. It tells us how B_{int} responds to permit price shocks over time, which governs the long run adjustment margin. Throughout, our identifying variation in permit prices comes from the carbon price surprise series.

4.1 Production Elasticities

We estimate the two production elasticities ξ_n and σ_n from firms' conditional factor demands. σ_n measures how easily firms substitute away from emissions at a given B_{int} , while ξ_n measures substitutability among capital, labor, and materials.

4.1.1 Inner Nest: Capital-Labor-Materials (ξ_n)

Under CES production, the optimal ratio of any two inputs j and k satisfies a relationship analogous to equation (6), governed by the common elasticity ξ_n . Our estimating equation for the materials-to-labor ratio is:

$$\log \frac{M_{int}}{L_{int}} = \xi_n \log P_{r,int}^L + \beta_M \log P_{int}^M + \lambda_{1,in} + \lambda_{1,ct} + \varepsilon_{1,int}, \quad (18)$$

where $\lambda_{1,in}$ is a firm fixed effect, $\lambda_{1,ct}$ is a country-year fixed effect that absorbs national wage and price trends, and $\varepsilon_{1,int}$ is the error term. The wage $P_{r,int}^L$ is our variable of interest and is endogenous. We label it with an additional subscript r because we will be using a regional wage consistent with our regional labor market instrumental variables approach.

The ξ_n coefficient on the log wage is the elasticity of substitution. When wages rise, firms substitute toward materials, and the extent of this substitution identifies ξ_n . The identifying variation we use is within-country-year variation in regional wages. We estimate the regression separately for each sector n and we two-way cluster standard errors by firm and year.²³ The CES structure implies that we can also estimate this elasticity from the capital-to-labor ratio, $\log(K/L)$, which yields the same ξ_n if the nested CES specification is correct. We estimate both and report the results in Appendix A.

The identification challenge is that wages are endogenous. Local wages may comove with unobserved shocks, such as productivity shocks, that also affect firms' input choices. We address this using a granular instrumental variable (GIV) strategy following Gabaix and Koijen (2024). The intuition is that since the firm size distribution in European labor markets is right-skewed, idiosyncratic labor demand shocks to large employers bid up local wages for all firms in the region, acting as residual labor supply shifters uncorrelated with any individual firm's productivity. We define labor markets at the region-year level and construct a leave-one-out instrument that aggregates idiosyncratic employment shocks to other firms, weighted by their employment shares. Technical details of the GIV construction are provided in Appendix F.

²³Under 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.

4.1.2 Outer Nest: Emissions-Composite (σ_n)

We estimate the outer nest elasticity σ_n between emissions and the productive composite using the analogous conditional factor demands in equation (5). Given the inner nest estimates ξ_n , we construct the composite price P_{int}^X as the normalized nested-CES price index over capital, labor, and materials. CES weights are given by baseline cost shares and substitution is governed by ξ_n . We measure the productive composite quantity as observed non-emissions spending deflated by the price index: $X_{int} = (C_{int}^K + C_{int}^L + C_{int}^M)/P_{int}^X$.

Our estimating equation is:

$$\log \frac{E_{int}}{X_{int}} = -\sigma_n \log P_t^E + \beta_X \log P_{int}^X + h_{2,n}(t) + \lambda_{2,in} + \varepsilon_{2,int}, \quad (19)$$

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 identify σ_n from the coefficient $-\sigma_n$ on the log emissions price. When permit prices rise, firms reduce emissions relative to the composite, and the magnitude of this response reveals σ_n .²⁴ The firm fixed effect absorbs the time-invariant share parameter and baseline emissions terms, and we include $\log P_{int}^X$ as a control with coefficient β_X that we do not interpret causally. We estimate this regression separately for each sector n and cluster standard errors by firm and year.²⁵ Our baseline specification uses a linear sector-specific trend.

Similar to the inner nest, the price of emissions is endogenous. As an example, when firms improve their clean technology, their production processes become less emissions-intensive, reducing aggregate permit demand and driving down equilibrium permit prices. The sector-specific time trends $h_{2,n}(t)$ absorb smooth technological progress, but concerns remain if technology and prices comove at higher frequencies. Identification is further complicated because all firms face a common nominal emissions price.

We instrument for the emissions price using the carbon price surprise series described in Section 2.2, aggregated to an annual frequency by summing within each calendar year. These high-frequency price movements around discrete regulatory announcements are allowance supply shocks, purged of slow-moving macroeconomic, financial, and energy market conditions, and plausibly uncorrelated with firms' technological choices, which evolve over longer horizons.²⁶

²⁴Under 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 the carbon price surprise series.

²⁵Because X_{int} and P_{int}^X depend on the first-stage estimate $\hat{\xi}_n$, standard errors come from a two-stage pairs cluster bootstrap that reestimates ξ_n in each replicate and propagates that uncertainty to σ_n . The firm and year variances are combined by two-way inclusion-exclusion (Cameron et al., 2011), with each bootstrap variance estimated from the interquartile range of the replicates, rescaled by the interquartile range of a standard normal, rather than from their standard deviation. The just-identified IV estimator has no finite moments (Kinal, 1980), so occasional extreme replicates—most likely when resampling the 17 year clusters—would dominate the standard deviation, while the quantiles of the replicate distribution remain well behaved. Appendix Figure A3 shows that inference is robust to a variety of other approaches. For the residual other non-manufacturing sector, we substitute parametric $\hat{\xi}_n$ draws from its asymptotic normal with dimension-matched cluster-robust standard errors, as the inner nest IV was numerically unstable on a subset of resamples.

²⁶Other emissions-related policies—such as renewable portfolio standard-style policies—do not threaten identification of σ_n as long as they are not time-varying or their time-varying component is unrelated to the allowance supply

Clean technology is what this regression leaves unexplained: the firm fixed effect, the sector trend, and the residual together capture the part of a firm’s emissions intensity that the permit and composite prices do not account for, which equation (5) identifies with $(\sigma_n - 1) \log B_{int}$. Appendix D shows how to invert the regression equation to recover B_{int} at the firm-year level.

Validation via Morishima Elasticities The nested CES restricts emissions to be equally substitutable with capital, labor, and materials. We test this assumption in Appendix A.1 by estimating pairwise Morishima elasticities between emissions and each productive input. Under nested CES all three should equal σ_n . We estimate Morishima elasticities that are uniformly small in magnitude and statistically indistinguishable from zero and across input pairs. An estimation approach allowing for a variable elasticity indicates they are relatively constant over the range of prices in the sample.

4.2 Reallocation Parameters

Carbon pricing raises unit costs more at dirty firms than at clean ones, shifting output shares across firms and across sectors. Three objects govern the aggregate elasticity of substitution summarizing this margin (Section 3.3): the emissions heterogeneity index χ_{nt} , the within-sector demand elasticity ϵ_n , and the cross-sector demand elasticity v . We recover each from the data in turn.

4.2.1 Heterogeneity Index χ_{nt}

We use equation (12) to calculate χ_{nt} for each sector-year directly from the data, evaluating each firm’s emissions cost share $s_{E,int}$ at the observed permit price and the composite prices P_{int}^X . Confidence intervals come from resampling firms within sector.

4.2.2 Demand Elasticities ϵ_n and v

Our model’s demand system (Section 3.2.2) is CES at both the firm and sector levels, with monopolistically competitive firms setting prices as a constant markup over marginal cost. The markup identity $\mu_n = \epsilon_n / (\epsilon_n - 1)$ in equation (7) lets us recover the within-sector demand elasticity ϵ_n from observed markups. We compute the sector-level markup μ_n by pooling revenue and variable costs across firms and years within each sector:

$$\mu_n = \frac{\sum_t \sum_{i \in \mathcal{I}_n} R_{int}}{\sum_t \sum_{i \in \mathcal{I}_n} (P_t^E E_{int} + P_{int}^X X_{int})},$$

where R_{int} is firm i ’s revenue. Given the CES structure, we then back out the demand elasticity as $\epsilon_n = \mu_n / (\mu_n - 1)$.

We estimate the cross-sector demand elasticity v from final demand, using nominal final demand expenditure by product from the Eurostat use tables and sector price indices from EU KLEMS

shocks in our instrument.

(Section 2.2).²⁷ Taking logs of the resulting demand curve gives our estimating equation:

$$\log V_{nct}^F = (1 - v) \log P_{nct} + \alpha_n + \alpha_{ct} + \varepsilon_{3,nct}, \quad (20)$$

where V_{nct}^F is nominal final demand expenditure on sector n in country c and year t , and P_{nct} is the sector's price index. We use expenditures instead of quantities because quantities are not reported in the data. To identify v we instrument $\log P_{nct}$ with the intermediate input price index. Shifts in this index move firms' materials costs and hence the prices firms charge households. Since households do not purchase intermediate inputs directly, the index affects final demand only through P_{nct} .

4.3 Clean Technology Dynamics

We take the dynamic extension of Section 3.4 to data by estimating equation (17) horizon by horizon using local projections. In our implementation, we replace clean technology with the observable composite-to-emissions ratio. In the case where $\sigma_n = 0$ as we will estimate below, using this observable ratio as the outcome rather than clean technology generates the same estimates, which tell us how much firms' clean technology adjusts at horizon h in response to a carbon price shock at $t - 1$ (see Appendix I). The estimating equation is:

$$\log \frac{X_{in,t+h}}{E_{in,t+h}} - \log \frac{X_{in,t-1}}{E_{in,t-1}} = \gamma_{h,n}^E \log P_{t-1}^E + \gamma_{h,n}^X \log P_{in,t-1}^X + \alpha_{in}^{(h)} + \eta_{in,t+h}, \quad h \in \{0, 1, 2, 3, 4\}. \quad (21)$$

The left-hand side cumulates changes between $t - 1$ and $t + h$, so the coefficient at $h = 4$ is a five-year cumulative elasticity. The dating of prices at $t - 1$ matches the time-to-build structure in equation (14). Prices at $t - 1$ shift clean technology investment $I_{in,t-1}^B$, which raises clean technology beginning in period t .

We estimate equation (21) separately by sector and compute the impulse response functions out to five years. Since the permit price is endogenous, we instrument for $\log P_{t-1}^E$ using the lagged carbon price surprise series as we did when estimating σ_n . All specifications include firm fixed effects. For inference, we use Driscoll-Kraay standard errors, which allow for serial correlation induced by the cumulative horizon- h outcome and for cross-sectional dependence. Appendix I shows that, under the implementation lag and no-depreciation assumptions, the local projection specification recovers a dynamic investment policy response to a change in the carbon price.

²⁷Three KLEMS industry groups each span two of our model sectors, so we merge those sector pairs and estimate on the resulting eight sector groups. Under CES demand with a common v , expenditure on any union of sectors follows the same demand curve with the union's CES price index, so the group-level regression identifies the same v . We exclude public administration, education, and health (NACE sections O, P, and Q), whose output is largely non-market and valued at cost.

5 Results

This section presents the estimates of the key model parameters, organized around the three margins developed in Section 3. Section 5.1 reports the firm-level production elasticities that identify the short-run input substitution margin. Section 5.2 reports the two ingredients that govern the reallocation margin: the emissions heterogeneity index χ_{nt} and the within- and cross-sector demand elasticities. Section 5.3 reports our local projection estimates of the long-run clean technology response to carbon pricing. Section 5.4 synthesizes these estimates into marginal abatement cost curves at the sector and economy-wide levels. Section 5.5 reports robustness checks.

5.1 Production Elasticities

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.

5.1.1 Inner Nest: Capital-Labor-Materials (ξ_n)

Figure A4 in Appendix A presents sector-specific estimates of ξ_n from equation (18). Estimates for most sectors lie near one, and for most sectors we cannot reject Cobb-Douglas ($\xi_n = 1$): point estimates lie modestly above one in some sectors (e.g., rubber and plastics) and below one in others (e.g., cement), but these differences are imprecisely estimated. The power and mining sectors have imprecise negative point estimates—neither rejects Cobb-Douglas—which we replace with the smallest positive sector estimate in quantification (Appendix A).²⁸ Appendix Figure A5 shows that these estimates are stable across different ways of constructing the wage GIV.

The inner nest elasticities affect our analysis through their role in constructing the composite input price P_{int}^X and quantity X_{int} . The variation in ξ_n across sectors implies different aggregation weights when combining capital, labor, and materials into the productive composite and shapes how firms respond to materials price fluctuations induced by permit price shocks.


5.1.2 Outer Nest: Emissions-Composite (σ_n)

The central parameter governing firm-level marginal abatement costs is the elasticity of substitution σ_n between emissions and the productive composite. Figure 2 presents our sector-specific estimates from equation (19), with OLS estimates in gray and IV estimates in black.

All estimates cluster near zero, with OLS and IV point estimates similar across all sectors. We reject Cobb-Douglas ($\sigma_n = 1$) for every sector, with all estimates economically close to, and statistically indistinguishable from, the Leontief case ($\sigma_n = 0$). The IV first stage is relatively strong: Kleibergen-Paap first-stage F statistics, reported in the right margin of Figure 2, exceed

²⁸Our main counterfactual results are robust to this imputation choice; for example, replacing both estimates with the median positive sector estimate instead leaves the aggregate counterfactual results essentially unchanged and preserves qualitative sector-level patterns.

Figure 2: Emissions-Composite Elasticity (σ_n)



Notes: This figure plots sector-specific estimates of the elasticity of substitution σ_n between emissions and the productive composite from equation (19). Gray circles show OLS estimates; black triangles show IV estimates using the carbon price surprise series. Error bars show 95% confidence intervals based on a two-stage cluster bootstrap by firm and year (see footnote 25 in Section 4.1.2). The dashed vertical line marks zero (Leontief). The right margin reports each sector's Kleibergen-Paap first-stage F statistic for the IV estimate.

17.7 in every sector. These uniform near-zero elasticities imply that firms have extremely limited ability to substitute away from emissions toward other inputs. The only way for firms to reduce emissions in the short run is to cut output.

The technological rigidity of the emissions input mix contrasts sharply with that of labor, capital, and materials in the inner nest. Firms have almost no ability to trade off the productive composite against emissions because emissions embodied in production are determined by the technology of the equipment already installed.

The finding that σ_n is approximately zero may appear counterintuitive for sectors like power, where the documented transition from coal to natural gas and renewables suggests substantial scope for input switching. The median firm in our panel operates only one facility (Figure A6), so the within-firm input substitution margin is mechanically narrow. Cost-weighting the regression does not reveal a positive elasticity masked at large firms (Appendix Table A2). In Appendix Section G we present case studies of the largest EU-ETS power firms to show why. These firms run baseload-dominant fleets, and as EU-ETS prices rose sharply over 2015–2021, coal capacity factors fell sharply at the firms with sizable coal and gas fleets while gas capacity factors at the same firms did *not* rise correspondingly. Coal-dominant firms, with no gas to switch to, kept running coal. In firms with matched generation data we see no strong evidence of within-firm input substitution. Baseload and flexible units are not substitutes within the firm even when the carbon price rises. The displaced coal generation flowed to renewables and imports at the system level, consistent with our overall finding that short-run emissions reductions in this setting operate primarily through across-firm reallocation rather than within-firm input substitution.

A further concern is that our outer nest elasticity estimates do not account for entry and exit of firms. If variation in carbon prices leads firms with different emissions intensities to enter or exit, our estimate—based purely on within-firm variation—might miss this margin of adjustment. Appendix Figure A7 presents an alternative set of outer nest elasticity estimates utilizing a random effects (at the firm level) estimator, which combines the within-firm variation used in our baseline with the between-firm variation from sources such as entry and exit driven compositional changes over time. The results suggest that accounting for entry and exit does not meaningfully affect the magnitude of the outer nest elasticity in our panel of EU firms.²⁹

5.2 Reallocation Parameters

We now turn to the reallocation margin, which operates across firms and sectors. The elasticity of sector-level marginal abatement curves depends not only on firm-level substitution elasticities but also on the scope for reallocation across heterogeneous firms. Two parameters govern the within-sector reallocation channel: the heterogeneity index χ_{nt} , which captures within-sector dispersion in emissions cost shares across firms, driven to a first order by clean technology, and the within-sector demand elasticity ϵ_n , which determines how readily consumers shift toward cleaner producers when carbon pricing raises dirty firms' costs. A third parameter, the cross-sector demand elasticity v , governs reallocation across sectors and the economy-wide elasticity. Figure 3 presents estimates of χ_{nt} , ϵ_n , and the sector-level elasticities of substitution; we report v separately below.

5.2.1 Emissions Heterogeneity Indices χ_{nt}

Figure 3a reveals substantial variation in emissions heterogeneity across sectors. The power sector exhibits the highest heterogeneity index at approximately 0.10, reflecting the wide dispersion in clean technology and emissions intensity across firms. At the other end, paper has an index of 0.01, indicating more homogeneous emissions profiles.³⁰ The magnitude of these heterogeneity indices matters for the aggregate elasticity. Since our estimated $\hat{\sigma}_n \approx 0$ for all sectors, the within-firm term vanishes and the sector-level elasticity is $\bar{\sigma}_{nt} = \chi_{nt}\epsilon_n$, adjusted in our quantification by the pass-through correction of Appendix H.2.


5.2.2 Within-Sector Demand Elasticities ϵ_n

Figure 3b reports the estimated demand elasticities, which vary by a factor of three across sectors. Refining exhibits the highest demand elasticity at 12, reflecting the commodity nature of refined fuels and consumers' ability to substitute across suppliers. Basic metals also shows a higher demand elasticity of 7. At the other extreme, rubber and plastics has a demand elasticity around 4,

²⁹An important caveat is that Orbis has been shown to be limited in its ability to recover aggregate patterns related to entry and exit (Bajgar et al., 2020; Kalemli-Özcan et al., 2024). So, while the estimates presented in Appendix Figure A7 suggest our results are robust to entry-exit patterns observed within Orbis, it is likely Orbis is not well suited to definitively capture the degree to which this channel can explain aggregate trends.

³⁰Appendix Figure A11 shows the underlying within-sector distribution of emissions cost shares, the key ingredient for the heterogeneity index, for the largest emitting sectors.

Figure 3: Heterogeneity Index, Demand Elasticities, and Aggregate Elasticities



Notes: Panel (a) shows the heterogeneity index χ_{nt} computed within each sector-year following equation (12). The plotted value is the average of χ_{nt} across sample years, denoted $\bar{\chi}_n$. Panel (b) shows sector demand elasticities ϵ_n estimated from observed markups: $\epsilon_n = \mu_n/(\mu_n - 1)$, where μ_n is the sector-level markup pooled across firms and years. Panel (c) shows the aggregate elasticity $\bar{\sigma}_n$ in the Leontief case, computed from the time-averaged heterogeneity index and demand elasticity using the exact expression derived in Appendix H.2. Error bars show 95% confidence intervals: panel (a) resamples firms within sector (999 replications), panel (b) uses the delta method, and panel (c) pairs each joint resample of the heterogeneity index and pass-through correction with a demand elasticity drawn from its asymptotic normal, reporting percentile intervals. The dashed line corresponds to the aggregate, economy-wide elasticity derived in Appendix H.3.

and power has a demand elasticity around 5, consistent with more differentiated products or less consumer flexibility.

5.2.3 Aggregate Elasticities by Sector $\bar{\sigma}_n$

Figure 3c shows the estimated aggregate elasticities implied by the demand elasticities and the heterogeneity index when we adopt the Leontief firm elasticity ($\sigma_n = 0$) found in Section 5.1.2. The top polluting sectors like power, refining, and basic metals have the most elastic response to emissions pricing. These sectors have substantial cross-firm heterogeneity in emissions intensity—especially power and basic metals—while refining and basic metals exhibit highly elastic demand.³¹

5.2.4 Cross-Sector Demand Elasticity v

We also estimate the elasticity of substitution across sectors, v , which governs reallocation of final demand from emissions-intensive to cleaner sectors. Table A4 in Appendix A reports estimates of equation (20). The OLS price coefficients in columns (1)–(3) exceed one, implying a negative v : demand shocks raise expenditure and prices together, biasing the coefficient upward.

To address this simultaneity, columns (4)–(6) instrument the price with the intermediate input price index, which yields a strong first stage. Our preferred specification is column (6), which uses total final use to measure final demand. Household and total final consumption in columns

³¹The confidence interval for fabricated metal products is asymmetric, with the point estimate near its upper limit. Dispersion in this sector comes from a small number of emissions-intensive firms whose emissions cost shares lie far above the sector median.

(4)–(5) give similar point estimates with wider confidence intervals. We estimate $v \approx 1.1$, close to Cobb-Douglas demand. Expenditure shares respond only mildly to relative prices and across-sector reallocation provides a modest additional margin for emissions reductions beyond within-sector reallocation. The aggregate economy-wide emissions elasticity using this demand elasticity estimate is plotted as the dashed line in Figure 3c. Its value is 0.3, comparable to the most elastic sectors.

5.3 Clean Technology Dynamics

Figure 4 plots our estimates of the five-year cumulative effect of a permit price shock on clean technology, $\hat{\gamma}_{4,n}^E$ from equation (21). The IV estimates span a modest range, approaching 0.6 for the most responsive sectors. Refining has the largest five-year point estimate ($\hat{\gamma}_4^E = 0.564$, SE 0.148), followed by power (0.425, SE 0.123) and mining and quarrying (0.256, SE 0.152). A positive elasticity means that a higher permit price raises the composite-to-emissions ratio: holding productive inputs fixed, these firms achieve lower direct emissions per unit of input, consistent with the directed technical change prediction that firms exposed to the carbon price invest in cleaner production. Because the carbon price surprise series isolates unanticipated movements in the permit price that are also likely perceived as persistent, we read these local projections as an approximation to the dynamic response of clean technology investment to a carbon price change.³²

The remaining sectors have smaller positive point estimates and fabricated metals is the only sector with a small negative point estimate ($\hat{\gamma}_4^E = -0.044$, SE 0.087), but it is statistically indistinguishable from zero. Chemicals (0.011, SE 0.064) and basic metals (0.033, SE 0.157) have essentially zero point estimates.³³ Appendix Figure A9 shows the full sector-level impulse response functions.


Estimating a pooled five-year elasticity yields $\hat{\gamma}_4^E = 0.211$ (SE 0.043) when firms are weighted equally, and 0.318 (SE 0.055)—about 50% larger—when weighting observations by its sector’s 2005 verified emissions share. The larger emissions-weighted response shows that carbon-intensive sectors—which account for the bulk of aggregate emissions—invest in clean technology more than the average firm. Appendix Figure A10 plots both pooled responses by horizon. In both, improvements accrue incrementally over the five years following the shock rather than on impact.

External Validation Against Colmer et al. (2024) Colmer et al. (2024) provide a useful benchmark for our estimates. Using a matched difference-in-differences design on a subset of French manufacturing, they estimate a treatment effect of -0.174 on logged emissions intensity during EU-ETS Phase II. Re-estimating our local projection on the corresponding manufacturing

³²Permit banking links current spot prices to future permit prices so that unanticipated price shocks signal persistent changes.

³³Investment that makes the composite input more productive, such as materials-saving efficiency improvements, would register as a *decline* in measured clean technology. In the presence of input-output linkages, a carbon price also incentivizes these kinds of investments. If this occurs in the data, our estimates understate the clean technology response.

Figure 4: EU-ETS Price Elasticity of Clean Technology



Notes: This figure plots sector-specific estimates of the five-year cumulative coefficient $\gamma_{4,n}^E$ from equation (21), whose outcome uses the composite quantity $X_{int} = (C_{int}^K + C_{int}^L + C_{int}^M)/P_{int}^X$. Gray circles show OLS estimates and black triangles show IV estimates. Each estimate comes from a separate sector-level local projection with firm fixed effects, and the baseline specification includes the $\log P_{in,t-1}^X$ control. The $t - 1$ permit price is instrumented with the lagged carbon price surprise series. Error bars show 95% confidence intervals based on Driscoll-Kraay standard errors. The dashed vertical line marks zero.

sectors yields $\hat{\gamma}_4^E = 0.119$ (standard error 0.067).³⁴ Because Colmer et al. (2024) estimate the treatment effect of introducing the ETS whereas we estimate a price elasticity, we need a baseline non-ETS carbon price faced by comparison firms to translate their effect into price units. Using baseline prices between 2 and 6.5 EUR/tCO₂—a range supported by French statutory fuel excises and OECD effective carbon-rate estimates—their treatment effect implies an elasticity of 0.086–0.157, which brackets our estimate (see Appendix J for details). Thus, despite using different samples and identification strategies, the two approaches imply similar longer-run responses to carbon prices by firms.


5.4 Marginal Abatement Cost Curves

We now construct marginal abatement cost (MAC) curves implied by our model. Figure 5 presents sector-level and economy-wide MAC curves on a log-log scale, averaged across the sample years as baseline-revenue-weighted geometric means. Each slope is inversely proportional to the corresponding aggregate elasticity as discussed in Section 3.3.

Panel (a) plots the sector MAC curves, with the four largest-emitting sectors highlighted: each

³⁴The overlapping sectors are basic metals; cement, lime, glass, and ceramics; chemicals; fabricated metal products; pulp, paper, and board; and rubber and plastic products. We also estimate a version including light manufacturing, which partially overlaps.

Figure 5: Sector Marginal Abatement Costs and the Role of Clean Technology



Notes: Panel (a) plots the sector aggregate MAC curves. The four largest-emitting sectors are highlighted and labeled, and the remaining sectors appear as thin unlabeled curves. Panel (b) plots the economy-wide aggregate MAC under baseline clean technology, the level prevailing in each sample year (solid blue), and under the advanced clean technology level induced by a doubling of the permit price (dashed green). The gray dashed horizontal lines mark the average 2005 permit price and twice that level. Both curves are marked at the doubled price, and the baseline-technology curve also at the 2005 price. All curves are displayed in log-log space. The sector and economy-wide aggregates are constructed by simulating our model on a grid of counterfactual permit prices in each year of 2005–2021, recovering the corresponding sector and economy-wide aggregate emissions, averaged across the sample years as baseline-revenue-weighted geometric means. Following the definition of the MAC, the relevant aggregate is held fixed along each curve: sector output for the sector curves and aggregate household final demand for the economy-wide curves.

traces how sector emissions fall as the carbon price rises, holding sector output fixed. Power—the largest emitter—also has the flattest curve, while the curves of several smaller sectors are close to vertical. The easiest abatement is concentrated precisely where emissions are largest. The median sector has an elasticity of about 0.2: by shifting output from dirtier to cleaner producers, a sector abates more cheaply than any of its constituent firms.

Aggregating one step further, the blue curve in Panel (b) plots the resulting economy-wide MAC at the baseline clean technology level prevailing over the sample years. It is flatter still than the median sector curve. Appendix Figure A12 isolates each margin’s contribution and shows that adding reallocation across sectors flattens the curve beyond what reallocation within sectors already delivers, but within-sector reallocation is the primary driver for why the economy-wide MAC is relatively flat.

Panel (b) also shows how improvements in clean technology shift the MAC downward. The dashed green curve traces the MAC schedule with clean technology improvements caused by a doubling of the carbon price, recovered using the five-year local projections estimate. Moving from the solid blue curve to the dashed green one lowers the MAC by about 14 USD/tCO₂ at the doubled price. Increasing the carbon price to twice its 2005 level delivers 53.0 million tons of additional abatement in total—36% more than the price increase would achieve by moving along the solid blue curve, holding clean technology fixed.³⁵

³⁵The elasticity of the economy-wide MAC is similar across the two curves since clean technology primarily shifts

5.5 Robustness

Appendix A collects robustness and specification checks for our two central estimation results: the near-zero outer nest elasticity σ_n and the clean technology response to carbon pricing. For the outer nest elasticity, the Morishima elasticities between emissions and each productive input (Section 4.1.2) are statistically indistinguishable from one another and uniformly close to zero, supporting the nested CES restriction that emissions are equally substitutable with capital, labor, and materials (Appendix Tables A10 and A11). The finding that $\sigma_n \approx 0$ is also robust to quadratic, cubic, and natural spline time trends, including under inference that remains valid when instruments are weak (Appendix Figures A13 and A15); to instrumenting with the carbon price surprise series and the wage GIV together (Appendix Figure A14); to the 2013 expansion of EU-ETS coverage under Phase III (Appendix Table A5); and to weighting firms by total input cost following Oberfield and Raval (2021) (Appendix Table A2), which provides no evidence that large firms have a meaningfully positive substitution elasticity. Nor is inference fragile to the small number of year clusters ($G = 17$): wild cluster bootstraps on the year dimension (Cameron et al., 2008; MacKinnon et al., 2021) yield similar confidence intervals (Appendix Figure A3). For the clean technology response, the estimated impulse response accumulates roughly linearly across horizons rather than arriving on impact or decaying, consistent with the one-period implementation lag and no-depreciation assumptions behind our law of motion (Appendix I). Inference for the local projections is likewise robust to wild cluster bootstrapping to handle the small number of year clusters (Appendix Figures A9 and A10).

6 Counterfactuals

We now use the estimated model to run counterfactuals that isolate the equilibrium effect of changes in the EU-ETS permit price. Our central counterfactual doubles the EU-ETS permit price—approximately its observed 2005–2021 increase—and traces the effect through the full production network. We distinguish a short run, in which clean technology is fixed, from a long run, in which firms’ clean technology investment responds over the following five years. We simulate this price shock in every sample year and report the baseline-revenue-weighted geometric mean effect across years.

6.1 Aggregate Results

Table 1 quantifies the effects of doubling the carbon price. The top row shows that, in the short run with clean technology fixed, doubling the carbon price cuts emissions by 23%. A higher cost of emissions results in output falling 3.6% and a corresponding increase in output prices of 2.8%. Since emissions and labor are complements, employment falls 1.0%.³⁶

the level of the curve rather than the slope.

³⁶For a marginal, 1% price change, Appendix Table A6 reports the corresponding elasticities: the emissions elasticity (-0.34) is an order of magnitude larger in absolute value than the output (0.04) and employment (0.01) elasticities.

Table 1: Equilibrium Effects of Doubling the Carbon Price

	Δ Emissions (%)		Δ Output (%)		Δ Prices (%)		Δ Labor (%)	
	Total	IO share	Total	IO share	Total	IO share	Total	IO share
<i>Reallocation: Active</i>								
Short run	-22.89%	0.07	-3.62%	0.83	2.80%	0.75	-0.98%	0.35
Long run	-31.15%	0.03	-2.51%	0.83	1.87%	0.74	-0.67%	0.41
<i>Reallocation: Suppressed</i>								
Short run	-4.70%	0.82	-3.87%	0.79	3.24%	0.74	-2.25%	0.63
Long run	-20.63%	0.12	-2.50%	0.78	2.07%	0.74	-1.49%	0.64

Notes: This table reports counterfactual outcomes for ETS-regulated firms from 2005 to 2021 under a doubling of the carbon price (the approximate 2005–2021 change). Each year’s equilibrium uses the firms observed in that year, and annual effects are baseline-revenue-weighted geometric means across years. *Reallocation: Active* allows within- and across-sector reallocation; *Reallocation: Suppressed* shuts down both margins. *Short run* holds clean technology fixed at baseline; *Long run* additionally incorporates the estimated, sectorally heterogeneous five-year change in clean technology. “IO share” is the share of each effect attributable to input-output network propagation, computed as the difference between the full-network equilibrium and an alternative equilibrium where the materials expenditure shares Ω_{mnt} in equation (8) are set to zero, expressed as a fraction of the total effect on the unit interval $[0, 1]$.

Over the long run, firms respond to the higher price by investing in clean technology. The second row includes the estimated five-year clean technology response and shows that emissions reductions deepen to 31%. Our counterfactual long-run prediction is broadly in line with the actual 35% decline in emissions at regulated installations over our sample, suggesting that carbon pricing was a central driver of the observed reduction. In addition to cutting emissions, clean technology eases other economic impacts in the long run. Output falls by only 2.5%, a 31% smaller fall than the short run, and the price and employment effects both shrink relative to the short run.

The lower two rows of Table 1 present results when suppressing reallocation by holding firms’ relative output shares fixed. Without reallocation, emissions fall by 5% in the short run, roughly a fifth of what is achieved with reallocation. Since firm production is Leontief, this emissions reduction comes from output curtailment. Reallocation dampens the effects on consumer prices and quantities because it allows consumers to shift from relatively high price to low price producers. It also significantly blunts the aggregate negative effect on employment as cleaner firms shrink less or even grow compared to a world without reallocation. This finding has consequences for policy design. Carbon markets often include provisions like production-indexed rebates or exemptions for trade-exposed sectors to reduce cost impacts of carbon pricing. Our results suggest that policies that limit cost impacts will mute reallocation and blunt precisely the margin that delivers the short run abatement.

The input-output share columns in Table 1 split each effect into a direct effect of the carbon price and an indirect effect transmitted through the production network. The emissions reduction is almost entirely direct: network propagation contributes just 1.5pp of the 23% total. The indirect effect is small for emissions because the shock the network transmits is far smaller and lands on the wrong firms. Sector materials prices rise by only about 2.9% versus a doubling of the carbon price. Emissions intensive firms also devote less of their costs to materials, 57% against a sector average

of 68%, with a within-sector correlation between carbon and materials cost shares of -0.45 . The indirect effect through the input-output actually lowers the relative price of emissions-intensive firms by 0.3pp rather than raising it.

Unlike emissions, the economic impacts are mostly indirect. Network propagation accounts for most of the output decline (83%), the price increase (75%), and about a third (35%) of the employment effect. Network propagation scales up the common level of marginal costs, which moves output, prices, and employment, but it generates almost none of the within-sector relative price dispersion that reallocates market share from dirty to clean firms and drives the emissions decline.³⁷

Appendix Table A7 presents results of the policy counterfactuals under different values of the demand elasticity, the production elasticities, and the choice of capital rental rate. The impacts on prices and output are relatively similar regardless of the parameter values. The impact on emissions, as one should expect, is sensitive to the demand elasticity and elasticity of substitution between emissions and the productive composite. More elastic demand and more elastic emissions substitution lead to greater emissions declines, as described above. The impact on labor is sensitive to the production elasticities. Greater ability to substitute in either nest leads to a smaller reduction in labor, or even an increase under some parameterizations.

6.2 Sector-Level Results


The aggregate effects mask wide variation across sectors. Figure 6 plots the sectoral impacts of doubling the carbon price, splitting each outcome into the short-run effect (blue) and the incremental five-year clean technology response (green). Their sum, labeled on each bar, is the long-run effect. Emissions fall the most in refining and in the power sector, which are the sectors with the greatest scope for reallocation given their elastic demand and within-sector dispersion in clean technology (Section 5.2.3). The clean technology contribution to emissions reduction is highly heterogeneous across sectors, reflecting the sector-level responses estimated in Section 5.3.

The price and output panels show where the economic impacts concentrate. Output prices rise in every sector, and by the most in the emissions-intensive ones such as cement, power, and basic metals, whose direct permit costs are largest and which bore the largest emissions declines. Output contracts most in these same sectors. Clean technology adoption blunts the price increases, again concentrated in power, and modestly cushions the output declines. Appendix Figure A16 decomposes the price changes and shows how the input-output network amplifies them.

Employment effects are the most uneven because of heterogeneity in the estimated inner nest elasticity ξ_n . Net job losses concentrate in the energy sectors—power, mining, and refining. Elsewhere, employment impacts are more mixed. As materials grow more expensive, firms substitute toward labor, so some sectors add workers even as their output falls. Clean technology adoption

³⁷One notable result is that the input-output share of the price effect is very similar whether reallocation is active or suppressed. To a first-order, price pass-through is governed by the materials shares Ω_{mnt} and so is invariant to quantity reallocation.

Figure 6: Sectoral Short-Run Impacts and the Effect of Long-Run Clean Technology




Notes: Panels (a)–(d) plot the change in emissions, labor employment, output prices, and output quantities by sector. The blue bar is the short-run effect of doubling the carbon price with clean technology fixed. The green bar is the additional effect of the estimated, sectorally heterogeneous five-year clean technology response, defined as the long-run effect minus the short-run effect. The sum is the long-run effect and is labeled on each bar. Annual sector changes are baseline-revenue-weighted geometric means across 2005–2021.

cushions the hardest-hit sectors, offsetting much of the short-run employment loss in refining and in power.

6.3 Decomposition: Role of Reallocation

Figure 7 decomposes each sector's emissions change into five channels, showing where the abatement originates. The channels correspond to a sequence of counterfactuals that switch on one adjustment margin at a time, so each bar is the incremental effect of a margin given the ones before it. Network propagation holds firm and sector market shares fixed and sector final demand quantities at baseline, so output moves only through demand transmitted along input-output linkages. The

Figure 7: Decomposition of Emissions Changes by Sector



Notes: This figure decomposes the change in emissions at ETS firms under a doubling of the carbon price and the associated estimated five-year change in clean technology. Network propagation captures the change in emissions from firms' output responses when firm and sector market shares remain fixed and sector final demand quantities are held at baseline levels, so demand shifts reach firms only through input-output linkages. Within-sector reallocation captures the additional change from shifting market shares across firms within sectors. Across-sector reallocation captures the additional change from shifting demand across sectors while holding aggregate final demand quantity fixed. Clean technology captures the additional change from the estimated sector-specific change in emissions productivity once all market shares adjust. Aggregate demand captures the additional emissions change when nominal spending is held fixed instead of real final demand quantity. Each cumulative effect is a baseline-revenue-weighted geometric mean across 2005–2021. The colored bars are differences between adjacent cumulative effects, are displayed in the order the effects are introduced, and sum exactly to the black percent change total.

two reallocation channels then let expenditure respond to the relative price changes induced by the carbon price, first across firms within sectors and next across sectors, holding the aggregate final demand quantity fixed. Clean technology then adds the estimated five-year change in clean technology with all market shares free to adjust, so it nets the direct decline in emissions per unit of output. Aggregate demand finally relaxes the fixed aggregate quantity, allowing demand to contract or grow at all firms.

Within-sector reallocation is the primary driver in most sectors and dominates in the highest-emitting sectors like refining, power, and cement. Across-sector reallocation is small everywhere: with $v \approx 1.1$, final demand is close to Cobb-Douglas, so relative price changes shift expenditure across sectors only weakly. Clean technology is one of the most heterogeneous channels, reflecting the sector-level responses estimated in Section 5.3. It has the largest impacts in refining and light manufacturing. Network propagation and aggregate demand contribute modest and fairly uniform emissions reductions. Network propagation is largest in sectors that sell a larger share of output

as materials to other firms, while the uniformity of the aggregate demand channel is because it is a common proportional contraction from an increase in the aggregate price index.

6.4 Robustness: Carbon Leakage to Unregulated Firms

A natural concern with the EU-ETS is that carbon pricing shifts production from regulated to unregulated firms, undoing some of the intended emissions reductions (Section 3.5). To quantify the extent of leakage we use our extended model with the larger Orbis panel of non-ETS firms (Section 2.3). Because emissions at non-ETS firms are unobserved, we impute emissions intensities using several different strategies. Appendix Table A9 reports leakage rates. In the long-run counterfactual with the estimated clean technology response, leakage rates across the five strategies range from -0.7 to 2.2%. Leakage rates can be negative if emissions-intensive non-ETS firms shrink. This can happen if the materials cost increase from propagation through input-output linkages is sufficiently large for non-ETS emitters, or if the clean technology improvement at ETS firms leads to substantial reallocation toward ETS firms. Equilibrium leakage to unregulated firms is limited and does not materially change our main conclusions, consistent with prior evidence (Naegele and Zaklan, 2019; Colmer et al., 2024).

7 Conclusion

This paper develops a framework for quantifying the equilibrium effects of carbon pricing using commonly available data on firm financials and emissions. We isolate three adjustment margins: firms' short-run input substitution, reallocation of output from dirty toward cleaner firms, and firms' long-run adoption of clean technology. We show how to aggregate firm-level responses into equilibrium outcomes at the sector and economy-wide levels. Importantly, our approach is portable: it can be applied wherever pollution pricing programs generate emissions data that can be merged with firm financial information, providing a template for evaluating pollution pricing policies across regulatory contexts.

Applying the framework to EU-ETS, we find the input substitution margin is rigid: the firm-level elasticity of substitution between emissions and other inputs is near zero in every sector, so firms cannot reduce emissions in the short run without cutting output. Reallocation across firms is a major driver of the aggregate emissions response, with the biggest reductions concentrated in sectors with substantial cross-firm heterogeneity in clean technology, such as power and cement, and in sectors facing especially elastic demand, such as refining. Lastly, we show that clean technology adoption responds to carbon price shocks, with the largest clean technology responses in carbon-intensive sectors.

These findings help reconcile a puzzle: how have European emissions fallen so sharply while output and employment have remained relatively stable? The answer is that reallocation across firms and endogenous clean technology adoption are the dominant adjustment margins, while short-run within-firm substitution plays almost no role. Large aggregate emissions reductions coexist with

moderate economic costs because market shares shift toward cleaner firms and because carbon pricing itself induces the technological progress that lowers emissions intensity over time. For policymakers, the implications are twofold. First, pollution pricing can deliver substantial emissions reductions without large economic disruptions, because abatement comes mainly from reallocation toward cleaner firms and from induced technology rather than from firms cutting output. Second, because clean technology takes years to build and pays off only if carbon prices stay high, the clean technology responses we estimate are consistent with firms investing when price increases are expected to persist. Credible, lasting policy is important for the long-run technology dividend.

Our main framework could be extended in several ways. Our headline counterfactuals are partial equilibrium, holding factor prices fixed, though we show in Appendix E that the approach can be easily extended to general equilibrium. The set of firms is fixed: entry, exit, and relocation abroad are absent, so our leakage estimates capture reallocation toward existing unregulated firms rather than the offshoring of production. Market structure is monopolistically competitive with constant markups, abstracting from strategic pricing responses to carbon costs and local concentration. And the clean technology response we estimate is a reduced-form deterministic investment margin of a set of potential pre-existing technologies rather than uncertain innovation.

References

- Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous (2012) “The Environment and Directed Technical Change,” *American Economic Review*, Vol. 102, No. 1, pp. 131–166.
- Atkeson, Andrew and Patrick J. Kehoe (1999) “Models of Energy Use: Putty-Putty versus Putty-Clay,” *American Economic Review*, Vol. 89, No. 4, pp. 1028–1043.
- Bajgar, Matej, Giuseppe Berlingieri, Sara Calligaris, Chiara Criscuolo, and Jonathan Timmis (2020) “Coverage and representativeness of Orbis data,” OECD Science, Technology and Industry Working Papers 6, OECD Publishing.
- Barrutiabengoa Ortubai, Joxe Mari, Rodrigo Enrique Falbo Piacentini, Agustín García Serrador, and Juan F. Rubio-Ramírez (2025) “Unraveling the impact of a carbon price shock on macroeconomic variables: a Narrative Sign Restrictions approach,” *BBVA Research Working Paper No. 25/04*.
- Baumol, William J. and Wallace E. Oates (1988) *The Theory of Environmental Policy*, Cambridge, UK: Cambridge University Press, 2nd edition.
- Bjørnland, Hilde Christiane, Jamie Cross, and Felix Kapfhammer (2024) *The drivers of emission reductions in the European carbon market*, No. 1/2024: Working Paper.
- 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.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller (2008) “Bootstrap-Based Improvements for Inference with Clustered Errors,” *Review of Economics and Statistics*, Vol. 90, No. 3, pp. 414–427.
- (2011) “Robust Inference with Multiway Clustering,” *Journal of Business & Economic Statistics*, Vol. 29, No. 2, pp. 238–249.
- Casey, Gregory (2024) “Energy efficiency and directed technical change: implications for climate change mitigation,” *Review of Economic Studies*, Vol. 91, No. 1, pp. 192–228.
- 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.
- Cropper, Maureen L. and Wallace E. Oates (1992) “Environmental Economics: A Survey,” *Journal of Economic Literature*, Vol. 30, No. 2, pp. 675–740.
- 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.
- Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer (2015) “The Next Generation of the Penn World Table,” *American Economic Review*, Vol. 105, No. 10, pp. 3150–3182.
- Gabaix, Xavier and Ralph SJ Koijen (2024) “Granular instrumental variables,” *Journal of Political Economy*, Vol. 132, No. 7, pp. 2274–2303.
- Hollingsworth, Alex, Carl Kitchens, Taylor Jaworski, and Ivan Rudik (2026) “Economic geography and air pollution regulation in the United States,” *Journal of Political Economy Microeconomics*, Vol. 4, No. 1, pp. 38–77.
- Jézéquel-Royer, Marie and Grégory Levieuge (2025) “Going Green: Will the R-Star Shine Brighter?,” https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5050773, SSRN Working Paper.
- Johansen, Leif (1959) “Substitution versus Fixed Production Coefficients in the Theory of Economic Growth: A Synthesis,” *Econometrica*, Vol. 27, No. 2, pp. 157–176.

- Kalemli-Özcan, Şebnem, Bent E. Sørensen, Carolina Villegas-Sanchez, Vadym Volosovych, and Sevcan Yeşiltaş (2024) “How to Construct Nationally Representative Firm-Level Data from the Orbis Global Database: New Facts on SMEs and Aggregate Implications for Industry Concentration,” *American Economic Journal: Macroeconomics*, Vol. 16, No. 2, p. 353–74.
- Känzig, Diego R. (2025) “The Unequal Economic Consequences of Carbon Pricing,” *NBER Working Paper No. 31221*.
- Kinal, Terrence W. (1980) “The Existence of Moments of k-Class Estimators,” *Econometrica*, Vol. 48, No. 1, pp. 241–249.
- 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].
- MacKinnon, James G., Morten Ørregaard Nielsen, and Matthew D. Webb (2021) “Wild Bootstrap and Asymptotic Inference with Multiway Clustering,” *Journal of Business & Economic Statistics*, Vol. 39, No. 2, pp. 505–519.
- 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.
- Martin, Ralf, Mirabelle Muûls, and Ulrich J. Wagner (2016) “The Impact of the European Union Emissions Trading Scheme on Regulated Firms: What Is the Evidence after Ten Years?” *Review of Environmental Economics and Policy*, Vol. 10, No. 1, pp. 129–148.
- 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.
- Naegele, Helene and Aleksandar Zaklan (2019) “Does the EU ETS cause carbon leakage in European manufacturing?” *Journal of Environmental Economics and Management*, Vol. 93, pp. 125–147.
- 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.
- Pethig, Rüdiger (1976) “Pollution, Welfare, and Environmental Policy in the Theory of Comparative Advantage,” *Journal of Environmental Economics and Management*, Vol. 2, No. 3, pp. 160–169.
- Petrick, Sebastian and Ulrich J Wagner (2014) “The impact of carbon trading on industry: Evidence from German manufacturing firms,” *Available at SSRN 2389800*.

- Popp, David (2002) “Induced Innovation and Energy Prices,” *American Economic Review*, Vol. 92, No. 1, pp. 160–180.
- 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.
- Siebert, Horst, Jürgen Eichberger, Ralf Gronych, and Rüdiger Pethig (1980) *Trade and Environment: A Theoretical Enquiry*, Amsterdam: Elsevier.
- World Bank (2026) “Carbon Pricing Dashboard,” <https://carbonpricingdashboard.worldbank.org/>, Accessed August 10, 2026.

Appendix Table of Contents

Appendix A	Additional Figures and Tables	A.2
A.1	Morishima Elasticities: Empirical Validation of Nested CES Specification	A.27
Appendix B	Data Construction and Measurement	A.32
Appendix C	Generalized Model of Production and Abatement	A.34
C.1	A General Emissions-as-an-Input Representation	A.34
C.2	A Separable Microfoundation	A.36
Appendix D	Model Quantification Details	A.41
D.1	Demand	A.41
D.2	Supply	A.42
D.3	Equilibrium Pricing and Revenue	A.43
D.4	Quantification of Remaining Model Parameters	A.43
D.5	Simulation of Counterfactuals	A.46
D.6	Emissions Imputation for Non-ETS Firms	A.49
Appendix E	General Equilibrium Extension	A.50
E.1	Income and Factor Market Equilibrium	A.50
E.2	General Equilibrium Solution Algorithm Sketch	A.51
Appendix F	Granular Instrumental Variable Construction	A.51
Appendix G	Input Substitution: An Electricity Market Case Study	A.54
Appendix H	Marginal Abatement Cost Derivations: Firm, Sector, and Economy	A.59
H.1	Firm-Level Marginal Abatement Costs	A.59
H.2	Sector-Level Marginal Abatement Costs	A.63
H.3	Economy-Wide Marginal Abatement Costs	A.69
Appendix I	Microfoundation for the Clean Technology Local Projection	A.84
I.1	Law of Motion with Time-to-Build Lag and No Depreciation	A.85
I.2	An Investment Rule as Directed Technical Change	A.85
I.3	Combined Structural Law of Motion	A.88
I.4	Moving Average Representation and Identification of σ_n	A.88
I.5	Local Projection Coefficient and Carbon Price Surprise Orthogonality	A.88
I.6	Mapping Clean-Technology Dynamics to Observables	A.90
Appendix J	External Benchmark: Translating the Colmer et al. (2024) Treatment Effect	A.91
J.1	Sector Coverage	A.91
J.2	Determining Baseline Prices P_0	A.91
Appendix K	Long Run Elasticity of Substitution with Endogenous Clean Technology	A.92

A Additional Figures and Tables

Figure A1 plots the emissions share by sector over our sample period. The emissions share of power has dropped by 28% since the beginning of EU-ETS, while the shares of basic metals and chemicals have increased by factors of 1.8 and 2.4, respectively. Patterns are similar when using the full EUTL sample of regulated installations.

Figure A2 graphs the intuition behind our estimation equations for the elasticities of substitution. The blue lines show the isoquants for two different production functions, one with isoquants that are approximately linear such that emissions and the composite are highly substitutable (σ'), and one with isoquants that are nearly kinked such that emissions and the composite are highly complementary (σ). The left panel plots the relative price of the two inputs as a downward sloping dotted line. Where the production function isoquants are tangent to this line denotes the optimal mix of emissions and composite inputs. For both isoquants and the initial relative price, this occurs at the same input bundle $(E_0, X_0) \equiv (E'_0, X'_0)$.

The right panel captures the effect of an increase in the emissions price, which steepens the relative price curve with the new relative price shown as a solid line. A cost-minimizing firm will reallocate inputs in response to this relative price change and the extent of reallocation will depend on the elasticity of substitution. If inputs are highly substitutable, the firm will substantially decrease emissions and increase the composite to (E'_1, X'_1) . If inputs are complementary, the extent of reallocation will be much more muted and the input mix barely changes to (E_1, X_1) . The response of relative inputs to exogenous changes in relative prices thus identifies the elasticity of substitution between the inputs.

Figure A3 investigates the robustness of inference for the outer nest production elasticity. Our baseline approach clusters standard errors by firm and year. We compare five standard error schemes. Three are two-stage cluster bootstrap schemes that each reestimate the inner nest $\hat{\xi}_n$ per replicate and combine block bootstrap variances following Cameron et al. (2011): clustering by firm, by firm and year (our maintained choice), and by firm and region-year. The remaining two are wild cluster bootstrap schemes on the year dimension (MacKinnon et al., 2021)—a year-only unrestricted wild bootstrap and a hybrid that adds the two-stage bootstrap firm variance to the wild bootstrap year variance—which address the small number of year clusters directly. Confidence intervals are not meaningfully different across the approaches except for firm-only clustering, which tends to produce smaller intervals.

Figure A4 plots our estimates of the inner nest elasticities from equation (18) which govern the substitutability between capital, labor, and materials. We find that the inner nest is generally more substitutable than the outer nest with meaningful variation across sectors. The materials-labor IV estimates span -3.9 to 1.6. Elasticity estimates near or below 1 are consistent with the prior literature documenting capital-labor complementarity (Oberfield and Raval, 2021). For most sectors, the OLS and IV estimates are similar. We estimate noisy negative elasticities for power and for mining and quarrying in the materials-labor specification, but neither is statistically distinguishable from Cobb-Douglas, and negative values cannot enter the CES composite. For the

simulations we impute both from the smallest positive sector estimate, which comes from cement. Estimates using the capital-labor ratio are also larger than the outer nest estimates and tend to cluster between 0 and 1. We test the robustness of our simulation results to instead using the capital-labor estimates in Table A7 below.

Figure A5 plots the inner nest elasticity estimates under different instruments and measures of wages, estimating the materials-labor first-order condition in panel (a) and the capital-labor condition in panel (b). The black circles are our main estimates reported in Figure A4. The blue triangles instead only omit the own firm from the construction of the instrument rather than the own sector, where the exclusion of data in the instrument is to avoid contamination by a firm's own idiosyncratic demand shocks. The gray squares do not omit any data from the instrument construction. The green stars return to the main instrument, which excludes a firm's own sector, but instead treat the endogenous regressor as the observed firm wage rather than the computed regional wage. Elasticity estimates are extremely robust to how we construct the instrument. Changing from the regional wage to the firm wage makes the estimates much noisier for some sectors—consistent with its weaker first stage, with a median Kleibergen-Paap F statistic across sectors of about 14 versus about 27 for the regional wage GIV specifications—but most are extremely similar to estimates from our main specification. The capital-labor condition behaves the same way: across the three regional wage GIV specifications the sector estimates in panel (b) never differ from our main capital-labor estimates by more than about 0.05, and the two sectors whose materials-labor estimates are negative, power and mining, have positive capital-labor point estimates.

Figure A6 displays sector-specific histograms of the distribution of ETS installation count across firms. ETS installations are the directly regulated units under the policy. An installation may be a single plant, but there can be multiple installations at the same site. The installation count therefore serves as an upper bound on the number of distinct plants or locations a firm operates and thus the scope for within-firm reallocation by shifting production across sites. The figure shows that the installation distribution is skewed strongly to the right. The vast majority of firms operate a single installation, and the largest sectoral average is only 2.8 installations per firm. There are a handful of firms operating over 10 installations, primarily in power and in mining.

Figure A7 examines how the outer nest elasticity estimates respond to allowing for firm entry and exit. Our baseline specification uses firm fixed effects, which identify the elasticity from within-firm variation over time. We compare these estimates to an otherwise-identical random effects specification that additionally exploits between-firm, cross-sectional variation. The gap between the two proxies for the scope for changes in the composition of firms—through entry and exit—to move the estimated elasticity. The two sets of estimates are broadly similar and remain economically small across sectors, although a Hausman test rejects the consistency of the random effects estimator in a handful of sectors, which we mark with an asterisk.

Figure A8 plots the sector-specific clean technology trends recovered from the estimated time trends in the outer nest specification. Aggregate clean technology increased by approximately 46% by 2021: the same emissions now yield 46% more output than in 2005. The aggregate improvement

masks sectoral heterogeneity. The power sector had the largest gains, reflecting the historical shift from coal to natural gas and renewable generation in the EU. Basic metals, chemicals, and mining show declining measured emissions productivity. Part of the chemicals decline reflects the Phase III (2013) expansion of EU-ETS coverage, which first brought chemicals' process emissions within the system. This raises measured emissions at the coverage boundary independent of any change in production technology. Note that these recovered trends are descriptive, not causal.

Figure A9 plots the full IV local projection impulse response functions for clean technology. We plot the main Driscoll-Kraay confidence intervals along with ones derived from standard errors from a year wild bootstrap. Sectors generally show an increasing trend in clean technology following a permit price shock, culminating in the five-year horizon estimates in Figure 4.

Figure A10 plots the pooled IV local projections for clean technology, both weighting firms equally and weighting sectors by their baseline emissions shares.

Figure A11 plots the distribution of emissions cost shares across sectors. Wide right tails in the highest-emitting sectors indicate substantial dispersion in carbon cost exposure, which is the source of reallocation-based emissions reductions in the model.

Figure A12 plots the economy-wide MAC curves under different reallocation restrictions. The green solid line is the baseline MAC curve. The blue dashed line is the MAC curve if reallocation across sectors is shut down, so there is only reallocation within sectors. The vertical red dotted line is the MAC curve if all cross-firm reallocation is shut down and abatement occurs within firms, which operate Leontief production technologies. Abatement thus only occurs through output curtailment because no reallocation, even through within-firm input substitution, is possible. Recall the MAC is the shadow cost of cutting emissions while holding the corresponding aggregate fixed. For the economy-wide curves here, the fixed aggregate is household final demand, the cost is the economy's primary resource cost, and the curves trace the implementation under marginal cost pricing (Appendix H.3). If all reallocation across firms is restricted, there is no reallocation and the MAC is vertical. Adding across-firm but within-sector reallocation leads to a dramatic flattening of the MAC. Additionally allowing for across-sector reallocation only has a modest additional effect on the MAC slope. Most of the reallocation occurs across firms within a sector, consistent with Figure 7.

Figure A13 plots our IV estimates of the outer nest composite-emissions elasticities under different time trends. The point estimates cluster near zero regardless of what degree polynomial time trend we use, or whether we use natural splines with five degrees of freedom for additional flexibility. The sector-level Kleibergen-Paap F statistics reported in the right margin of each panel remain strong under the linear trend. The more flexible trends absorb substantially more of the identifying time-series variation, and the weakest first stage falls to 4.2. For those specifications we can rely on the Anderson-Rubin weak-IV confidence sets shown in Figure A15, which remain tight around zero.

Figure A14 plots our IV estimates of the outer nest composite-emissions elasticities under different instrumental variables approaches. All three estimates instead treat the relative price of

the composite to emissions as the endogenous variable of interest, rather than separating them and using the composite price as a control. The black circles are estimates using the carbon price surprise series as the instrument. The gray triangles are estimates that instead instrument using the wage GIV we use to identify the inner elasticity. The blue squares are estimates that use both instruments. The wage GIV estimates are generally noisy because the first stage is weak. The median Kleibergen-Paap F statistic for this instrument is near 1. Although the GIV is relatively strong for wages, it is not quite as strong for the computed price of the composite or emissions. The composite is a non-linear function of wages as well as other prices, and emissions are only indirectly affected by wages through equilibrium forces. The point estimates for the specifications using the carbon price surprise series IV generate results close to zero similar to our main specification which only instruments for the permit price.

Figure A15 plots IV estimates of the outer nest composite-emissions elasticities under different specifications, using Anderson-Rubin weak IV inference following recommendations in Andrews et al. (2019). The black circle plots IV estimates from our main specification with linear time trends, the dark gray triangle is quadratic time trends, the gray square is cubic time trends, and the blue diamond keeps linear time trends but treats the relative price of the composite to emissions as the endogenous regressor as in Figure A14. The legend reports each specification's median Kleibergen-Paap first-stage F statistic across sectors. The thick shaded bars are the weak IV confidence sets. Across all specifications, the weak IV confidence sets are relatively tight.

Figure A16 shows the change in sectoral prices when doubling the emissions price, including the associated change in clean technology computed using estimates from Figure 4. The blue part of the bar shows the direct effect of the permit price and clean technology, while the amber bar shows the additional effect from the shock being amplified by the input-output network. In all sectors, input-output propagation of the shock accounts for a substantial share of the price increase.

Table A1 provides a summary of our data coverage by sector. In key polluting sectors like power and cement we observe hundreds of regulated firms.

Table A2 presents estimates of our outer nest composite-emissions elasticity when weighting observations by firm costs following Oberfield and Raval (2021). Larger firms in terms of costs likely operate more installations (e.g., Figure A6), affording greater ability for within-firm, but across-installation reallocation in response to permit price shocks. A cost-weighted regression will put greater weight on these larger firms with more reallocation opportunities. We find that cost-weighted elasticities are generally close to zero and similar to our main, unweighted estimates, except for refining, which we find to be statistically significantly negative when cost weighted.

Table A3 presents the estimates corresponding to Figure A9.

Table A4 presents the estimates of our cross-sector demand elasticity described in the main text. Column (6) is our preferred specification: total final use as the outcome, instrumenting the sector price with the intermediate input price index, with sector and country-by-year fixed effects. Oberfield and Raval (2021) estimate the analogous cross-industry demand elasticity with a supply-side cost instrument and find estimates centered around one, in line with our preferred estimate of

1.1.

Table A5 tests the robustness of our composite-emissions elasticity to the 2013 expansion of EU-ETS coverage under Phase III. In the table we exclude 2013 from the sample, add a post-2013 break to the sector time trend, and estimate separate pre- and post-2013 elasticities within one specification. All three alternatives leave every sector’s estimate small or negative, far from the Cobb-Douglas value of one.

Table A6 reports carbon price elasticities corresponding to Table 1, computed from a 1% increase in the allowance price. Each entry gives the percent change in the outcome per 1% increase in the permit price.


Table A7 presents a set of robustness exercises for our policy counterfactual simulations with reallocation. Panel A varies the inner demand elasticity governing substitution across firms within a sector. The demand elasticity estimate is particularly important for emissions since it determines the ability of the economy to reallocate away from dirty firms. Panel B varies the outer cross-sector demand elasticity governing across-sector reallocation. The simulation results are all relatively unaffected by its value. Panel C varies the inner nest production elasticity of substitution between capital, labor, and materials. The value of this elasticity matters for employment outcomes. Permit price shocks propagate into materials prices and lead to substitution between materials and labor, making employment sensitive to the value of the inner nest production elasticity. Panel D varies the outer nest production elasticity of substitution between the productive composite and emissions. As expected, greater elasticities lead to steeper emissions responses and blunted impacts on prices and output. Panel E varies how we compute the rental rate of capital by using only subsets of the full Jorgenson formula, all omitting the capital gains/losses channel present in the main text. The rental rate computation matters for price and output. Adding a risk premium to the base real rate increases rental rates and blunts the impact of permit price shocks on output and firm prices. The baseline rental rate further accounts for depreciation and investment price growth and additionally blunts impacts.

Table A8 sorts firms into quartiles of emissions cost share π_{in} and reports how baseline emissions and revenue are distributed across those quartiles. The highest quartile holds 62% of emissions but only 9% of revenue. In the fixed-share scale exercise, output and emissions decline by the same common proportion in every quartile, so the table documents baseline concentration rather than heterogeneous contraction across firms.

Table A9 presents our leakage estimates using the 3,410,623 distinct firms in the full Orbis sample—not just firms in ETS—under alternative strategies for imputing emissions at non-ETS firms. Section D.6 describes the matching procedure. Without endogenous clean technology, leakage rates across the five imputation strategies range from 1.3% to 6.1%. With clean technology, leakage rates range from -0.7% to 2.2%. These results align with Colmer et al. (2024) and Naegele and Zaklan (2019), who find limited or no leakage under EU-ETS.³⁸


³⁸Colmer et al. (2024) find no evidence for leakage in shifting of input suppliers, product market reallocation, or within-firm reallocation across regulated and unregulated plants. Naegele and Zaklan (2019) find no evidence of changes in embodied carbon of imports. Our leakage analysis captures product market reallocation across firms and

Figure A1: Emissions Share by Sector Over Time




Notes: This figure plots each sector's share of total verified emissions in our EUTL-Orbis matched sample. The five sectors with the largest emissions shares in 2021 are labeled.

Figure A2: Elasticity of Substitution Between Emissions and Composite Input



Notes: This figure illustrates input responses to an emissions price increase for firms with different production technologies. Both isoquants produce the same output and are tangent to the initial relative price line at the same input bundle. The isoquant labeled σ has low substitutability between the composite input (X) and emissions (E); the isoquant labeled σ' has higher substitutability and lies weakly inside σ away from the tangency. Panel (a) shows the common optimal input choice (X_0, E_0 and X'_0, E'_0) under the initial relative price (dotted gray line). Panel (b) shows responses (X_1, E_1 and X'_1, E'_1) to an emissions price increase; each black line is the new isocost line tangent to one isoquant, and the dotted gray line reproduces the initial relative price line from Panel (a).


Figure A3: Outer Nest Clustering and Wild Bootstrap Robustness



Notes: This figure shows the maintained IV emissions-composite elasticity under five inference approaches. Point estimates are identical across schemes by construction. Schemes marked with a † reestimate the inner nest $\hat{\xi}_n$ in every bootstrap replicate, propagating first-stage uncertainty into σ_n , and combine block bootstrap variances by two-way inclusion-exclusion (Cameron et al., 2011) with a robust interquartile scale. Gray circles (†): firm clusters. Black diamonds (†): firm and year, the scheme reported throughout the paper. Blue squares (†): firm and region-year. Orange triangles: year-only unrestricted wild cluster bootstrap (WCU), ξ_n held fixed (MacKinnon et al., 2021). Purple stars: hybrid firm and year scheme—the two-stage bootstrap firm variance added to the WCU year variance, omitting the intersection correction, which is slightly conservative. Wild bootstrap schemes do not propagate inner nest uncertainty, so they omit the two-stage estimation uncertainty that the maintained scheme incorporates. Error bars show 95% confidence intervals. The dashed vertical line marks zero. The carbon price surprise instrument varies only across 17 years, so the wild bootstrap schemes address the small number of year clusters (Cameron et al., 2008).


within-firm reallocation and thus provides a structural complement to Colmer et al. (2024).

Figure A4: 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 (18). Panel (a) estimates the materials-labor first-order condition; Panel (b) estimates the capital-labor first-order condition. Gray circles show OLS estimates; black triangles show IV estimates using granular instrumental variables for wages following Gabaix and Koijen (2024). Unlike the matched ETS sample used elsewhere, these regressions use the broader Orbis GIV panel spanning 1998–2021. The 11 sector regressions contain 25.8 million observations; Appendix F describes the construction. Error bars show 95% confidence intervals from standard errors clustered by firm and year. Confidence intervals extending beyond the $[-5, 5]$ axis range are truncated at the panel edge. The right margin of each panel reports the sector's Kleibergen-Paap first-stage F statistic for the IV estimate. The dashed vertical lines mark zero (Leontief) and one (Cobb-Douglas).

Figure A5: Capital-Labor-Materials Elasticity of Substitution Robustness




Notes: This figure plots sector-specific estimates of the inner nest elasticity of substitution ξ_n among capital, labor, and materials from equation (18). Panel (a) estimates the materials-labor first-order condition; Panel (b) estimates the capital-labor first-order condition. Black circles show our main IV where we exclude a firm's own SIC 3 sector from the construction of the instrument; blue triangles show estimates when constructing the instrument only excluding the firm; gray squares show estimates when constructing the instrument without any exclusion of data; green stars show estimates when using the observed firm wage instead of the computed regional wage as the endogenous regressor. Legend labels report each specification's median Kleibergen-Paap first-stage F statistic across sectors, computed within the panel. Error bars show 95% confidence intervals from standard errors clustered by firm and year. Confidence intervals extending beyond the $[-5, 5]$ axis range are truncated at the panel edge. The dashed vertical lines mark zero (Leontief) and one (Cobb-Douglas).

Figure A6: EU-ETS Installations by Firm




Notes: This figure shows the distribution of EU-ETS installations by firm in our sample of firms used to simulate the model. The last bin captures all firms with 10 or more installations. The number reported in each panel is the sector's average installation count.

Figure A7: Outer Nest Entry-Exit Robustness




Notes: This figure shows sector-specific estimates of the outer nest emissions-composite elasticity of substitution σ_n from equation (19), comparing the maintained firm fixed effects specification against an otherwise-identical firm random effects specification. Both specifications instrument the log emissions price with the carbon price surprise series. The dependent variable is the log of the emissions-to-composite ratio. Blue triangles show the baseline fixed effects IV estimates, which coincide with the maintained estimates in Figure 2. Red circles show the random effects IV estimates. The fixed effects estimator identifies σ_n from within-firm variation over time, whereas the random effects estimator additionally exploits between-firm, cross-sectional variation. Error bars show 95% confidence intervals based on standard errors clustered by firm and year. Sectors marked with an asterisk are those for which a Hausman test rejects the consistency of the random effects estimator at the 5% level. The dashed vertical line marks zero.

Figure A8: Recovered Clean Technology Trends




Notes: This figure plots the sector-specific clean technology trends recovered from the linear time trends $h_{2,n}(t)$ in the IV specification of equation (19). Each curve is normalized to 1 in 2005. The thick line shows the aggregate trend from a single regression pooling all sectors. The dashed horizontal line marks the 2005 normalization of 1.

Figure A9: EU-ETS Clean Technology Local Projection Impulse Responses by Sector




Notes: This figure plots the sector-specific local projection impulse response functions from equation (21), whose outcome is the cumulative change in $\log(X/E)$. The baseline specification includes the $\log P_{in,t-1}^X$ control. Each estimate comes from a separate sector-level local projection with firm fixed effects. The $t-1$ permit price is instrumented with the lagged carbon price surprise series. Each point estimate carries two 95% confidence intervals. Blue error bars use Driscoll-Kraay standard errors, and red error bars use a wild cluster bootstrap on the year dimension (Webb weights, 9,999 replications). The dashed horizontal line marks zero.

Figure A10: Pooled EU-ETS Clean Technology Local Projection Impulse Responses: Firm-Average and Emissions-Weighted




Notes: This figure plots the pooled local projection impulse response function for the composite-to-emissions ratio $\log(X/E)$ from equation (21). The baseline specification includes the $\log P_{in,t-1}^X$ control. Each horizon is a separate pooled local projection with firm fixed effects, instrumenting the $t - 1$ permit price with the lagged carbon price surprise series. In the left panel, every firm-year is weighted equally. In the right panel, observations are weighted so that each sector's total weight equals its share of 2005 (EU-ETS Phase I baseline) emissions, with firms weighted equally within sector. Each point estimate carries two 95% confidence intervals. Blue error bars use Driscoll-Kraay standard errors, and red error bars use a wild cluster bootstrap on the year dimension (Webb weights, 9,999 replications). The dashed horizontal line marks zero.

Figure A11: Distribution of Emissions Cost Shares Across Sectors



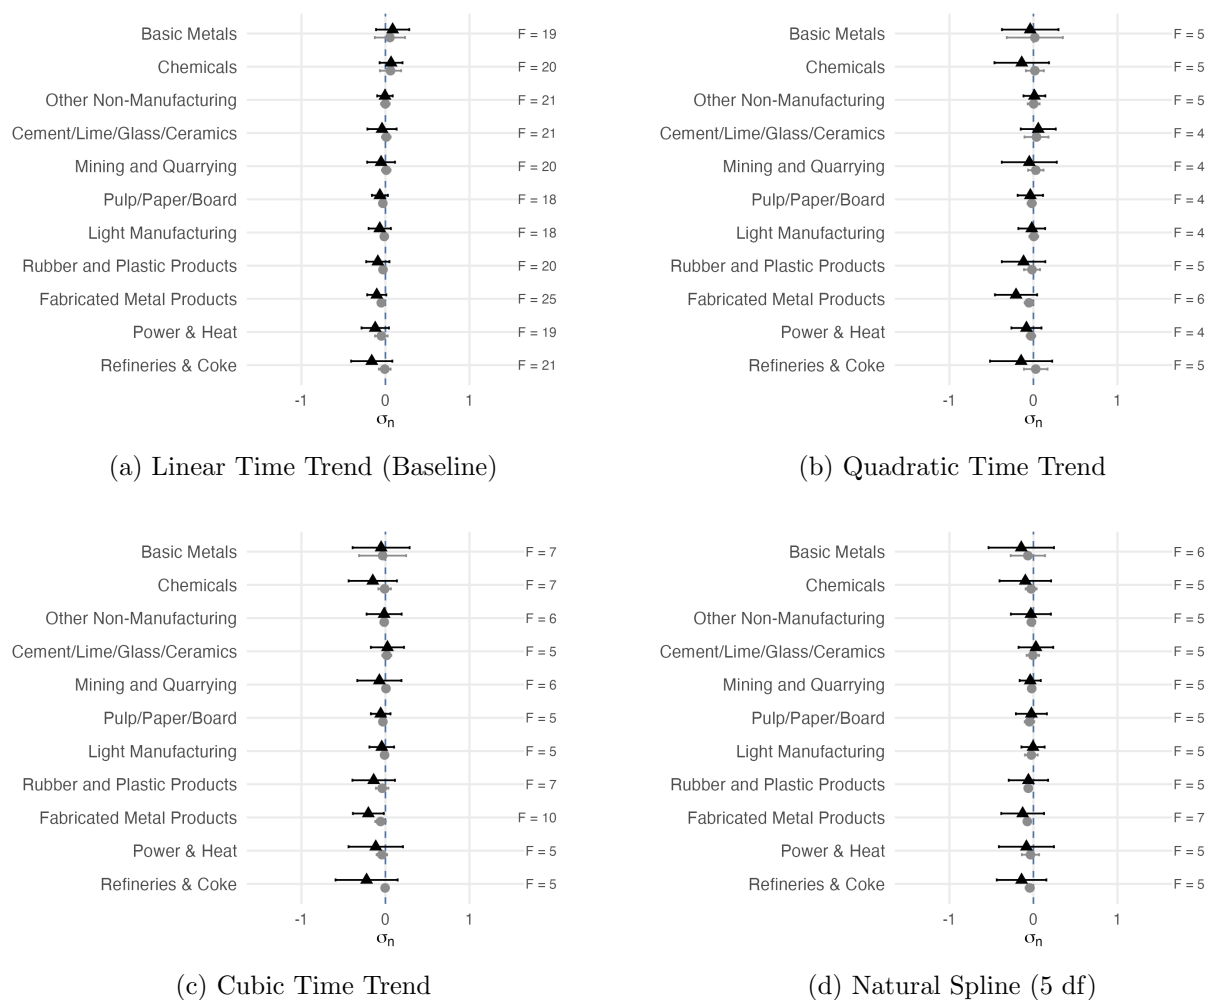
Notes: This figure shows the 2005 distribution of emissions cost shares across firms in all 11 model sectors. Each ridge depicts the within-sector distribution of emissions costs—tons of CO₂ times the average permit price—relative to firms' total input cost (the sum of capital, labor, materials, and emissions costs). Red diamonds mark sector medians. The horizontal axis is truncated, so the widest ridges extend beyond the displayed range.

Figure A12: Economy-Wide MAC Curves Under Alternative Reallocation Mechanisms




Notes: This figure plots the economy-wide planner MAC curve, constructed as the baseline-revenue-weighted geometric mean of year-specific curves across 2005–2021, under different reallocation mechanisms. Aggregate household final demand is held fixed along each curve. All curves hold clean technology fixed. The red dotted line incorporates only within-firm adjustment. The blue dashed line incorporates reallocation across firms within sectors. The green solid line incorporates full reallocation and matches the economy-wide MAC curve in Figure 5.

Figure A13: Emissions-Composite Elasticity Robustness to Time Trend Specification




Notes: This figure shows emissions-composite elasticity estimates under alternative time trend specifications. Gray circles show OLS estimates. Black triangles show IV estimates using the carbon price surprise series. Error bars show 95% confidence intervals based on a two-stage cluster bootstrap by firm and year (see footnote 25 in Section 4.1.2). The dashed vertical line in each panel marks zero (Leontief). The right margin of each panel reports the sector-level Kleibergen-Paap first-stage F statistics for the IV estimates.

Figure A14: Outer Nest Alternative Instruments




Notes: This figure shows emissions-composite elasticity estimates under alternative sets of instruments when instrumenting for the price of the composite relative to the price of emissions instead of treating the composite price as a control. All specifications use the baseline linear time trends. The black circles are estimates using the carbon price surprise series. The gray triangles use the wage granular instrumental variable instead of the carbon price surprise series. The blue squares include both the carbon price surprise series and wage granular instrumental variables. Legend labels report each instrument set's median Kleibergen-Paap first-stage F statistic across sectors. Error bars show 95% confidence intervals based on a two-stage cluster bootstrap by firm and year (see footnote 25 in Section 4.1.2). The dashed vertical lines mark zero (Leontief) and one (Cobb-Douglas). Off-scale point estimates are placed at the relevant boundary and labeled with their actual values, while confidence intervals are truncated at the boundary.

Figure A15: Outer Nest Weak IV Robustness



Notes: This figure shows emissions-composite elasticity estimates under alternative specifications and reports weak IV confidence sets. The black circle is our main specification with linear time trends, the dark gray triangle uses quadratic time trends, the gray square uses cubic time trends, and the blue diamond keeps the linear time trends but instruments for the relative emissions composite price instead of controlling for the composite price. Legend labels report each specification's median Kleibergen-Paap first-stage F statistic across sectors. Thin capped lines show 95% confidence intervals based on a two-stage cluster bootstrap by firm and year (see footnote 25 in Section 4.1.2). Teal outlined bands show Anderson-Rubin weak IV confidence sets following Andrews et al. (2019). The dashed vertical line marks zero (Leontief).

Figure A16: Sectoral IO Price Effects



Notes: This figure plots sectoral price changes from doubling the emissions price and adding the estimated five-year change in clean technology. The blue segment is the direct effect when input-output propagation is shut down. The amber segment is the additional effect transmitted through the input-output network. The two segments sum to the total sectoral price change.

Table A1: Sector Data Coverage Summary

Sector	NACE Codes	# Observations	# Unique Firms	Share of Emissions (%)
Power & Heat	35	8,332	1,039	46.44
Cement/Lime/Glass/Ceramics	23	8,270	1,027	15.21
Light Manufacturing	10–16, 18, 21, 26–33	6,489	806	2.33
Pulp/Paper/Board	17	4,410	485	2.91
Other Non-Manufacturing	All remaining NACE codes	3,267	527	5.52
Chemicals	20	3,162	362	6.74
Basic Metals	24	2,178	259	7.12
Mining and Quarrying	05–09	617	63	2.73
Refineries & Coke	19	617	62	10.19
Rubber and Plastic Products	22	491	61	0.18
Fabricated Metal Products	25	453	69	0.62

Notes: This table summarizes sectoral coverage in the matched 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.

Table A2: Outer Nest σ_n : Robustness to Cost Weighting

Sector	Unweighted		Cost-weighted		N
	OLS $\hat{\sigma}_n$	IV $\hat{\sigma}_n$	OLS $\hat{\sigma}_n$	IV $\hat{\sigma}_n$	
Basic Metals	0.053 (0.092)	0.086 (0.101)	0.078 (0.133)	0.124 (0.149)	2,178
Cement/Lime/Glass/Ceramics	0.010 (0.026)	-0.041 (0.089)	0.006 (0.020)	0.027 (0.054)	8,270
Chemicals	0.061 (0.064)	0.067 (0.070)	0.044 (0.112)	0.033 (0.152)	3,162
Fabricated Metal Products	-0.050 (0.013)	-0.103 (0.058)	0.001 (0.033)	-0.008 (0.032)	453
Light Manufacturing	-0.013 (0.017)	-0.067 (0.068)	0.005 (0.042)	0.036 (0.059)	6,489
Mining and Quarrying	0.010 (0.026)	-0.052 (0.085)	-0.004 (0.060)	-0.390 (0.297)	617
Other Non-Manufacturing	-0.001 (0.027)	-0.006 (0.047)	0.004 (0.034)	-0.026 (0.099)	3,267
Power & Heat	-0.048 (0.038)	-0.121 (0.083)	0.026 (0.049)	-0.051 (0.090)	8,332
Pulp/Paper/Board	-0.031 (0.015)	-0.066 (0.049)	-0.009 (0.017)	-0.009 (0.065)	4,410
Refineries & Coke	-0.008 (0.036)	-0.163 (0.125)	-0.035 (0.047)	-0.277 (0.106)	617
Rubber and Plastic Products	-0.029 (0.022)	-0.090 (0.071)	-0.027 (0.017)	-0.147 (0.114)	491

Notes: This table reports sector-level estimates of the outer nest elasticity σ_n across all four combinations of identification strategy (OLS vs. IV) and sample weighting (unweighted vs. cost-weighted). Cost weights are total input cost $C_K + C_L + C_M + C_E$ following Oberfield and Raval (2021). Cost weights are normalized to be mean one within each sector to stabilize the weighted demeaning. The IV columns instrument the log emissions price with the carbon price surprise series. Standard errors, from a two-stage cluster bootstrap by firm and year, are in parentheses.

Table A3: Sector-Level Local Projection IV Estimates

Sector	$\hat{\gamma}_0^E$	$\hat{\gamma}_1^E$	$\hat{\gamma}_2^E$	$\hat{\gamma}_3^E$	$\hat{\gamma}_4^E$	$N_{h=4}$
Basic Metals	0.095 (0.079)	0.339 (0.153)	0.421 (0.209)	0.197 (0.189)	0.033 (0.157)	922
Cement/Lime/Glass/Ceramics	0.046 (0.046)	0.115 (0.068)	0.147 (0.095)	0.189 (0.095)	0.139 (0.083)	3,416
Chemicals	0.128 (0.069)	0.325 (0.115)	0.245 (0.165)	0.119 (0.088)	0.011 (0.064)	1,401
Fabricated Metal Products	0.011 (0.033)	0.134 (0.171)	0.158 (0.235)	0.197 (0.126)	-0.044 (0.087)	150
Light Manufacturing	0.074 (0.036)	0.136 (0.063)	0.206 (0.061)	0.236 (0.083)	0.248 (0.070)	2,741
Mining and Quarrying	0.016 (0.080)	0.253 (0.159)	0.478 (0.195)	0.476 (0.162)	0.256 (0.152)	351
Other Non-Manufacturing	-0.031 (0.076)	0.058 (0.057)	0.075 (0.093)	0.039 (0.085)	0.109 (0.137)	1,128
Power & Heat	-0.043 (0.067)	-0.014 (0.129)	0.270 (0.107)	0.451 (0.112)	0.425 (0.123)	3,239
Pulp/Paper/Board	0.035 (0.024)	0.151 (0.053)	0.122 (0.059)	0.125 (0.075)	0.165 (0.075)	2,075
Refineries & Coke	0.140 (0.092)	0.375 (0.136)	0.609 (0.238)	0.491 (0.180)	0.564 (0.148)	294
Rubber and Plastic Products	0.099 (0.051)	0.224 (0.104)	0.379 (0.131)	0.366 (0.107)	0.199 (0.072)	207

Notes: This table reports sector-by-sector IV estimates of the cumulative permit price elasticity $\hat{\gamma}_{h,n}^E$ at horizons $h \in \{0, 1, 2, 3, 4\}$ from equation (21). The outcome is the cumulative change in $\log(X/E)$. The $t - 1$ permit price is instrumented with the lagged carbon price surprise series. $N_{h=4}$ reports the regression sample at the five-year horizon. Driscoll-Kraay standard errors are in parentheses.

Table A4: Cross-Sector Demand Elasticity Estimates

	Log Expenditure					
	(1)	(2)	(3)	(4)	(5)	(6)
Log(P)	1.721 (0.7864)	1.712 (0.7920)	1.338 (0.5442)	-0.2551 (0.7243)	-0.2609 (0.7367)	-0.1114 (0.3552)
Outcome	Household Cons.	Total Final Cons.	Total Final Use	Household Cons.	Total Final Cons.	Total Final Use
Model	OLS	OLS	OLS	IV	IV	IV
Instrument	–	–	–	II Price	II Price	II Price
Observations	1,722	1,722	1,734	1,722	1,722	1,734
Kleibergen-Paap F	–	–	–	83	82	86
Sector Group Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports the coefficient on log price, $1 - v$, from the regression $\log V_{nct}^F = (1 - v) \log P_{nct} + \alpha_n + \alpha_{ct} + \varepsilon_{nct}$ in equation (20), where V_{nct}^F is nominal final demand expenditure on sector n , P_{nct} is the sector's price index, and c indexes countries. The outcome is household final consumption in columns (1) and (4), total final consumption in columns (2) and (5), and total final use in columns (3) and (6). Columns (1)–(3) report OLS estimates. Columns (4)–(6) instrument for price using the intermediate input price index. Column (6) is our preferred specification. The Kleibergen-Paap row reports the first-stage F statistic for the excluded instrument in the IV columns. Standard errors, two-way clustered by country and year, are in parentheses.

Table A5: Outer Nest σ_n : Robustness to the 2013 EU-ETS Coverage Expansion

Sector	Baseline	Excl. 2013	2013 break	Period-specific $\hat{\sigma}_n$		N
				Pre-2013	Post-2013	
Basic Metals	0.086 (0.101)	0.085 (0.150)	−0.068 (0.100)	0.151 (0.117)	−0.144 (0.149)	2,178
Cement/Lime/Glass/Ceramics	−0.041 (0.089)	−0.048 (0.125)	−0.124 (0.076)	0.019 (0.077)	−0.181 (0.073)	8,270
Chemicals	0.067 (0.070)	0.067 (0.114)	−0.020 (0.077)	0.133 (0.101)	−0.090 (0.118)	3,162
Fabricated Metal Products	−0.103 (0.058)	−0.090 (0.060)	−0.139 (0.082)	−0.091 (0.069)	−0.166 (0.083)	453
Light Manufacturing	−0.067 (0.068)	−0.056 (0.090)	−0.124 (0.069)	−0.031 (0.074)	−0.159 (0.049)	6,489
Mining and Quarrying	−0.052 (0.085)	−0.047 (0.097)	−0.116 (0.063)	−0.002 (0.101)	−0.145 (0.082)	617
Other Non-Manufacturing	−0.006 (0.047)	0.006 (0.087)	−0.041 (0.086)	0.028 (0.073)	−0.075 (0.095)	3,267
Power & Heat	−0.121 (0.083)	−0.081 (0.090)	−0.136 (0.106)	−0.095 (0.094)	−0.163 (0.097)	8,332
Pulp/Paper/Board	−0.066 (0.049)	−0.059 (0.065)	−0.102 (0.061)	−0.030 (0.056)	−0.136 (0.046)	4,410
Refineries & Coke	−0.163 (0.125)	−0.088 (0.143)	−0.269 (0.167)	−0.101 (0.180)	−0.300 (0.142)	617
Rubber and Plastic Products	−0.090 (0.071)	−0.075 (0.102)	−0.178 (0.104)	−0.043 (0.097)	−0.246 (0.085)	491

Notes: This table reports the IV emissions-composite elasticity $\hat{\sigma}_n$. *Baseline* reproduces the IV estimates in Figure 2. *Excl. 2013* drops the year in which EU-ETS coverage expanded to these sectors' process emissions under Phase III. *2013 break* adds a post-2013 indicator to the sector time trend. *Pre-2013* and *Post-2013* come from a single specification that adds an interaction between the instrumented permit price and the post-2013 indicator to the baseline, so the two columns report the elasticity before and after the coverage expansion. The permit price is instrumented with the carbon price surprise series. Standard errors, from a two-stage cluster bootstrap by firm and year, are in parentheses in all columns (see footnote 25 in Section 4.1.2 and Figure A3).

Table A6: Carbon Price Elasticities

	ε_P		ε_Y		ε_E		ε_L	
	Total	IO sh.	Total	IO sh.	Total	IO sh.	Total	IO sh.
<i>Reallocation: Active</i>								
Short run	0.033	75%	−0.045	81%	−0.344	5%	−0.013	27%
Long run	0.023	74%	−0.032	82%	−0.525	3%	−0.009	37%
<i>Reallocation: Suppressed</i>								
Short run	0.033	75%	−0.041	80%	−0.050	83%	−0.024	65%
Long run	0.023	74%	−0.029	79%	−0.333	11%	−0.017	65%

Notes: This table reports elasticities of aggregate outcomes with respect to the EU-ETS permit price, computed from a 1% price increase. Each entry gives the percent change in the outcome per 1% increase in the permit price. Row definitions follow Table 1. The columns report the same four outcomes in elasticity form, ordered as prices (ε_P), output (ε_Y), emissions (ε_E), and labor (ε_L). “IO sh.” is the share of each effect attributable to input-output network propagation, expressed as a percentage of the total effect rather than the fraction reported in Table 1.

Table A7: Policy Counterfactuals Robustness

	ΔP (%)	ΔY (%)	ΔE (%)	ΔL (%)
<i>Panel A: Within-Sector Demand Elasticity</i>				
Double CO ₂ price				
Demand elasticity $\times 0.5$	3.15 (+2.34)	-3.50 (-2.76)	-15.11 (-1.64)	-1.25 (-0.45)
Demand elasticity $\times 2$	2.42 (+1.82)	-4.00 (-3.52)	-34.07 (-1.70)	-0.61 (-0.18)
Demand elasticity = estimated min	2.88 (+2.15)	-3.58 (-2.94)	-20.60 (-1.75)	-0.97 (-0.38)
Demand elasticity = estimated max	2.35 (+1.76)	-4.11 (-3.66)	-36.26 (-1.90)	-0.42 (-0.15)
Double CO ₂ price + clean tech				
Demand elasticity $\times 0.5$	2.05 (+1.50)	-2.34 (-1.83)	-26.33 (-0.98)	-0.85 (-0.33)
Demand elasticity $\times 2$	1.67 (+1.24)	-2.85 (-2.50)	-38.78 (-1.15)	-0.40 (-0.16)
Demand elasticity = estimated min	1.91 (+1.41)	-2.44 (-1.99)	-29.80 (-1.16)	-0.67 (-0.29)
Demand elasticity = estimated max	1.63 (+1.22)	-2.96 (-2.62)	-40.38 (-1.31)	-0.27 (-0.17)
<i>Panel B: Across-Sector Demand Elasticity</i>				
Double CO ₂ price				
Across-sector demand elasticity $\times 0.5$	2.81 (+2.10)	-3.63 (-3.02)	-22.61 (-1.68)	-1.07 (-0.43)
Across-sector demand elasticity $\times 2$	2.79 (+2.09)	-3.59 (-2.99)	-23.45 (-1.13)	-0.80 (-0.18)
Double CO ₂ price + clean tech				
Across-sector demand elasticity $\times 0.5$	1.87 (+1.39)	-2.52 (-2.09)	-31.01 (-1.20)	-0.71 (-0.33)
Across-sector demand elasticity $\times 2$	1.87 (+1.39)	-2.49 (-2.06)	-31.44 (-0.84)	-0.60 (-0.17)
<i>Panel C: Inner Nest Production Elasticity</i>				
Double CO ₂ price				
Inner nest elasticity $\times 0.5$	2.98 (+2.29)	-3.50 (-2.90)	-22.68 (-1.31)	-1.37 (-0.71)
Inner nest elasticity $\times 2$	2.38 (+1.66)	-3.83 (-3.22)	-23.38 (-1.93)	-0.48 (+0.11)
Inner nest elasticity = estimated min	3.05 (+2.36)	-3.40 (-2.80)	-22.58 (-1.20)	-1.68 (-1.00)
Inner nest elasticity = estimated max	2.34 (+1.61)	-3.52 (-2.90)	-23.39 (-1.79)	-0.20 (+0.41)
Inner nest elasticity = capital-labor estimates	2.43 (+1.69)	-3.26 (-2.63)	-23.29 (-1.54)	-1.03 (-0.32)
Double CO ₂ price + clean tech				
Inner nest elasticity $\times 0.5$	1.99 (+1.51)	-2.45 (-2.02)	-31.03 (-0.99)	-0.97 (-0.54)
Inner nest elasticity $\times 2$	1.59 (+1.10)	-2.63 (-2.19)	-31.45 (-1.31)	-0.28 (+0.07)
Inner nest elasticity = estimated min	2.04 (+1.56)	-2.38 (-1.95)	-30.97 (-0.93)	-1.18 (-0.74)
Inner nest elasticity = estimated max	1.57 (+1.07)	-2.43 (-1.99)	-31.43 (-1.11)	-0.14 (+0.25)
Inner nest elasticity = capital-labor estimates	1.64 (+1.14)	-2.26 (-1.81)	-31.35 (-0.80)	-0.75 (-0.29)
<i>Panel D: Outer Nest Production Elasticity</i>				
Double CO ₂ price				
Outer nest elasticity = 0.5	2.44 (+1.83)	-2.88 (-2.35)	-41.74 (-0.20)	-0.31 (-0.00)
Outer nest elasticity = 1.5	1.86 (+1.40)	-1.76 (-1.35)	-67.75 (+0.62)	0.65 (+0.49)
Outer nest elasticity = 2.5	1.43 (+1.09)	-1.00 (-0.68)	-82.71 (+0.57)	1.23 (+0.79)
Double CO ₂ price + clean tech				
Outer nest elasticity = 0.5	1.68 (+1.25)	-2.06 (-1.67)	-44.01 (-0.18)	-0.19 (-0.03)
Outer nest elasticity = 1.5	1.37 (+1.02)	-1.33 (-1.02)	-63.15 (+0.64)	0.54 (+0.36)
Outer nest elasticity = 2.5	1.13 (+0.85)	-0.81 (-0.55)	-75.81 (+0.81)	1.03 (+0.62)
<i>Panel E: Capital Rental Rate</i>				
Double CO ₂ price				
Real return	4.26 (+3.48)	-5.35 (-4.70)	-25.66 (-2.33)	-1.01 (-0.32)
Real return + risk premium	3.33 (+2.61)	-4.23 (-3.61)	-23.73 (-1.73)	-0.98 (-0.30)
Double CO ₂ price + clean tech				
Real return	2.82 (+2.28)	-3.69 (-3.22)	-32.81 (-1.61)	-0.71 (-0.27)
Real return + risk premium	2.22 (+1.72)	-2.94 (-2.49)	-31.68 (-1.23)	-0.68 (-0.25)

Notes: This table reports robustness of counterfactual outcomes to different parameter values. Reported results are baseline-revenue-weighted geometric means across years. Values in parentheses are the percentage-point contribution of the input-output network, calculated as the full-network outcome change minus the outcome change when the materials shares Ω_{mnt} are set to zero. Panel A varies the within-sector demand elasticity. Panel B varies the cross-sector demand elasticity. Panel C varies the inner nest elasticity across capital, labor, and materials. Panel D varies the outer nest elasticity between emissions and the productive composite. Panel E varies the capital rental rate calculation.

Table A8: Emissions and Revenue Shares by Emissions Cost Share

Quartile	Mean π_{in} (%)	Emissions Share (%)	Revenue Share (%)
Q1 (Lowest)	0.06	5.9	52.1
Q2	0.46	11.4	23.3
Q3	1.82	20.8	15.2
Q4 (Highest)	12.00	61.9	9.4
All Firms	1.51	100.0	100.0

Notes: This table sorts firms into quartiles of the baseline emissions cost share π_{in} , re-forming the quartiles within each year, and reports each quartile's shares of baseline emissions and revenue. Emissions cost share is the ratio of emissions costs to total input cost (the sum of capital, labor, materials, and emissions costs). The associated fixed-share scale exercise holds all within-sector and across-sector market shares at baseline levels, giving every quartile the same proportional output and emissions change. Those duplicate change columns are omitted because they do not measure heterogeneous contraction. Entries are averages weighted by baseline revenue across 2005–2021.

Table A9: Carbon Leakage to Unregulated Firms

Scenario	Nearest-neighbor imputation				Sector mean Revenue-weighted
	Inputs	Cost ratios	Revenue	All covariates	
Short Run	1.98 %	1.33 %	3.87 %	3.46 %	6.13 %
Long Run	-0.45 %	-0.73 %	0.61 %	0.20 %	2.22 %

Notes: This table reports leakage rates from ETS to non-ETS firms in the full Orbis sample under five emissions intensity imputation strategies. Leakage is $-100 \times \Delta E_{\text{non-ETS}} / \Delta E_{\text{ETS}}$, so a positive value means non-ETS emissions rise as ETS emissions fall. Appendix D.6 describes the matching procedure. The first four columns use nearest-neighbor imputations based on input levels, input cost-to-revenue ratios, revenue, or all three. The final column uses the revenue-weighted mean emissions intensity of ETS firms in the same country, sector, and year.

A.1 Morishima Elasticities: Empirical Validation of Nested CES Specification

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. Here we test the nested CES functional form itself by estimating 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:

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

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 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 $M_{KE} \neq M_{LE} \neq M_{ME}$.

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, similar to our outer nest specification in equation (19):

$$\log \frac{j_{ft}}{E_{ft}} = M_{jE} \log P_t^E + \sum_{k \neq E} \beta_{jk} \log P_{k,t} + \gamma_f + h_n(t) + u_{ft} \quad (22)$$

where f indexes firms, n indexes sectors, and t indexes years. We include firm fixed effects γ_f and linear time trends $h_n(t)$, matching the baseline outer nest specification. We instrument for the endogenous emissions price using the carbon price surprise series. We control for all three factor prices.

Table A10 reports OLS and IV estimates of the three Morishima elasticities. The third column for each elasticity additionally controls for log sales as a proxy for holding output fixed as the definition requires.³⁹ All three are uniformly small in magnitude, with a maximum absolute IV estimate of 0.050. All three are statistically indistinguishable from each other. Emissions appear to be similarly complementary with all productive inputs, consistent with treating capital, labor, and materials as a composite in the outer nest.

Figure A17 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.

Tables A11–A14 present quadratic IV estimates of the Morishima elasticities to test whether the elasticity of substitution is constant over the range of permit prices in the data. In these

³⁹By not holding output fixed, we are implicitly imposing homotheticity.

specifications we instrument for the linear and quadratic permit prices with the linear and quadratic surprise series. The tables report the coefficient estimates as well as the predicted elasticity at the minimum, median, and maximum annual average permit price. The first table shows pooled estimates. None of the estimates or predicted elasticities are statistically distinguishable from zero. The next three tables show sector-specific estimates. The tables report the minimum across the two first stages of the F statistic that both excluded instruments are jointly zero. In many cases this minimum first-stage F statistic is low, so we interpret these results with caution; however, only 4 of 99 predicted elasticities are statistically different from zero.


The evidence presented here informs our nested CES specification: Morishima elasticities are uniformly low, statistically indistinguishable across input pairs, and remain close to zero across the entire support of permit prices in the data. These findings support a representation of capital, labor, and materials as a CES composite in the CES outer nest along with emissions.

Table A10: Morishima Elasticities of Substitution

	$\log\left(\frac{K}{E}\right)$			$\log\left(\frac{L}{E}\right)$			$\log\left(\frac{M}{E}\right)$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\log(P^E)$	-0.0040 (0.0227)	-0.0180 (0.0618)	-0.0208 (0.0712)	-0.0029 (0.0144)	-0.0045 (0.0464)	-0.0047 (0.0501)	0.0025 (0.0223)	-0.0498 (0.0688)	-0.0284 (0.0480)
Model	OLS	IV	IV	OLS	IV	IV	OLS	IV	IV
Observations	38,286	38,286	38,131	38,286	38,286	38,131	38,286	38,286	38,131
Kleibergen-Paap F	–	17	17	–	17	17	–	17	17
Firm Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Input Price Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sectoral Linear Trends	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Log Sales Control	No	No	Yes	No	No	Yes	No	No	Yes

Notes: This table reports pooled OLS and IV estimates of the Morishima elasticities M_{jE} between each productive input and emissions from equation (22). All regressions control for labor, capital, and materials prices, firm fixed effects, and a sector linear time trend. Columns (3), (6), and (9) additionally control for log firm sales. IV regressions instrument for $\log(P^E)$ using the carbon price surprise series. Standard errors, from a cluster bootstrap by firm and year, are in parentheses.

Figure A17: Morishima Elasticities by Sector



Notes: Each panel shows sector-level estimates of the Morishima elasticity of substitution between the indicated input and emissions, with sectors ordered by the IV estimate. Gray circles are OLS estimates. Black triangles are IV estimates using the carbon price surprise series. Error bars show 95% confidence intervals based on a cluster bootstrap by firm and year. The dashed vertical line in each panel marks zero. All regressions control for labor, capital, and materials prices, firm fixed effects, and a sector linear time trend. Kleibergen-Paap first-stage F statistics for the IV estimates are shown in the right margin of each panel.

Table A11: Quadratic IV Morishima Elasticities

	$\log\left(\frac{K}{E}\right)$ (1)	$\log\left(\frac{L}{E}\right)$ (2)	$\log\left(\frac{M}{E}\right)$ (3)
$\log(P^E)$	0.7955 (0.8259)	0.4611 (0.5713)	1.0813 (1.0602)
$\log(P^E)^2$	-0.1338 (0.1326)	-0.0766 (0.0936)	-0.1861 (0.1738)
Observations	38,286	38,286	38,286
Min. First-Stage F	9	9	9
Elasticity at Minimum Price	0.80 (0.83)	0.46 (0.57)	1.09 (1.06)
Elasticity at Median Price	0.02 (0.09)	0.02 (0.06)	-0.00 (0.09)
Elasticity at Maximum Price	-0.31 (0.35)	-0.17 (0.26)	-0.46 (0.44)
Firm Fixed Effects	Yes	Yes	Yes
Input Price Controls	Yes	Yes	Yes
Sectoral Linear Trends	Yes	Yes	Yes

Notes: This table reports pooled quadratic IV estimates of the Morishima elasticities M_{jE} between each productive input and emissions from equation (22). All regressions control for labor, capital, and materials prices, firm fixed effects, and a sector linear time trend. Regressions instrument for $\log(P^E)$ and $\log(P^E)^2$ using the carbon price surprise series and its square. Standard errors, from a cluster bootstrap by firm and year, are in parentheses.

Table A12: Quadratic IV Morishima Elasticities: Capital-Emissions

	Basic Metals	Cement/Lime/Glass/Ceramics	Chemicals	Fabricated Metals	Light Manufacturing	Mining and Quarrying	Other Non-Mfg.	Power & Heat	Pulp/Paper/Board	Refineries & Coke	Rubber & Plastics
$\log P^E$	0.46 (0.89)	1.45 (1.61)	-1.94 (1.23)	2.13 (1.20)	0.10 (0.51)	1.70 (0.80)	1.06 (1.37)	1.27 (1.09)	0.39 (0.60)	0.87 (1.25)	2.09 (1.80)
$(\log P^E)^2$	-0.06 (0.15)	-0.26 (0.26)	0.35 (0.22)	-0.38 (0.22)	-0.01 (0.09)	-0.30 (0.17)	-0.17 (0.24)	-0.21 (0.18)	-0.06 (0.11)	-0.16 (0.24)	-0.35 (0.29)
Min. first-stage F	8.9	9.1	10.1	10.1	8.0	9.1	8.6	8.0	8.3	12.1	11.1
Elasticity at min. price	0.46 (0.89)	1.46 (1.62)	-1.95 (1.23)	2.14 (1.21)	0.10 (0.51)	1.70 (0.81)	1.06 (1.38)	1.27 (1.09)	0.39 (0.61)	0.87 (1.26)	2.10 (1.80)
Elasticity at median price	0.12 (0.12)	-0.05 (0.15)	0.11 (0.11)	-0.09 (0.04)	0.03 (0.06)	-0.03 (0.11)	0.05 (0.11)	0.05 (0.11)	0.05 (0.07)	-0.05 (0.16)	0.02 (0.12)
Elasticity at max. price	-0.02 (0.43)	-0.69 (0.59)	0.99 (0.62)	-1.04 (0.51)	0.00 (0.23)	-0.76 (0.57)	-0.37 (0.59)	-0.48 (0.48)	-0.10 (0.28)	-0.44 (0.61)	-0.85 (0.67)
Observations	2,178	8,270	3,162	453	6,489	617	3,267	8,332	4,410	617	491

Notes: This table reports sector-specific capital-emissions quadratic IV estimates of the Morishima elasticity M_{KE} between capital and emissions from equation (22). All regressions control for labor, capital, and materials prices, firm fixed effects, and a sector linear time trend. Regressions instrument for $\log(P^E)$ and $\log(P^E)^2$ using the carbon price surprise series and its square. Standard errors, from a cluster bootstrap by firm and year, are in parentheses.

Table A13: Quadratic IV Morishima Elasticities: Labor-Emissions

	Basic Metals	Cement/Lime/Glass/Ceramics	Chemicals	Fabricated Metals	Light Manufacturing	Mining and Quarrying	Other Non-Mfg.	Power & Heat	Pulp/Paper/Board	Refineries & Coke	Rubber & Plastics
$\log P^E$	-0.18 (0.69)	0.74 (0.72)	-1.67 (1.06)	1.07 (0.84)	0.27 (0.47)	0.79 (0.42)	0.27 (0.64)	0.69 (0.61)	0.62 (0.63)	1.30 (1.77)	0.59 (0.59)
$(\log P^E)^2$	0.04 (0.13)	-0.13 (0.12)	0.30 (0.18)	-0.18 (0.03)	-0.04 (0.08)	-0.13 (0.07)	-0.04 (0.10)	-0.11 (0.10)	-0.11 (0.11)	-0.23 (0.31)	-0.10 (0.11)
Min. first-stage F	8.9	9.1	10.1	10.1	8.0	9.1	8.6	8.0	8.3	12.1	11.1
Elasticity at min. price	-0.18 (0.69)	0.74 (0.73)	-1.68 (1.07)	1.07 (0.84)	0.27 (0.48)	0.79 (0.42)	0.27 (0.65)	0.70 (0.61)	0.62 (0.63)	1.31 (1.78)	0.60 (0.59)
Elasticity at median price	0.06 (0.08)	-0.01 (0.09)	0.09 (0.10)	-0.00 (0.06)	0.01 (0.05)	0.03 (0.06)	0.05 (0.06)	0.02 (0.11)	-0.00 (0.05)	-0.05 (0.16)	-0.00 (0.05)
Elasticity at max. price	0.17 (0.38)	-0.33 (0.32)	0.84 (0.44)	-0.45 (0.22)	-0.10 (0.22)	-0.30 (0.27)	-0.05 (0.27)	-0.26 (0.32)	-0.27 (0.26)	-0.62 (0.81)	-0.25 (0.26)
Observations	2,178	8,270	3,162	453	6,489	617	3,267	8,332	4,410	617	491

Notes: This table reports sector-specific labor-emissions quadratic IV estimates of the Morishima elasticity M_{LE} between labor and emissions from equation (22). All regressions control for labor, capital, and materials prices, firm fixed effects, and a sector linear time trend. Regressions instrument for $\log(P^E)$ and $\log(P^E)^2$ using the carbon price surprise series and its square. Standard errors, from a cluster bootstrap by firm and year, are in parentheses.

Table A14: Quadratic IV Morishima Elasticities: Materials-Emissions

	Basic Metals	Cement/Lime/Glass/Ceramics	Chemicals	Fabricated Metals	Light Manufacturing	Mining and Quarrying	Other Non-Mfg.	Power & Heat	Pulp/Paper/Board	Refineries & Coke	Rubber & Plastics
$\log P^E$	-0.32 (1.05)	0.18 (0.48)	0.83 (1.10)	1.13 (1.14)	1.16 (1.01)	1.14 (0.72)	0.49 (0.83)	2.54 (2.04)	0.56 (0.57)	2.02 (2.60)	1.71 (1.54)
$(\log P^E)^2$	0.06 (0.17)	-0.03 (0.09)	-0.14 (0.18)	-0.20 (0.18)	-0.20 (0.18)	-0.21 (0.13)	-0.08 (0.16)	-0.43 (0.34)	-0.11 (0.11)	-0.36 (0.45)	-0.29 (0.28)
Min. first-stage F	8.9	9.1	10.1	10.1	8.0	9.1	8.6	8.0	8.3	12.1	11.1
Elasticity at min. price	-0.32 (1.05)	0.18 (0.49)	0.84 (1.11)	1.14 (1.15)	1.17 (1.02)	1.15 (0.72)	0.49 (0.84)	2.55 (2.05)	0.56 (0.58)	2.03 (2.61)	1.72 (1.54)
Elasticity at median price	0.06 (0.12)	0.02 (0.07)	0.04 (0.09)	-0.01 (0.09)	-0.03 (0.08)	-0.09 (0.14)	0.03 (0.07)	0.06 (0.15)	-0.08 (0.07)	-0.08 (0.15)	0.03 (0.15)
Elasticity at max. price	0.22 (0.44)	-0.04 (0.27)	-0.30 (0.46)	-0.50 (0.49)	-0.54 (0.47)	-0.62 (0.42)	-0.17 (0.44)	-1.00 (0.77)	-0.35 (0.31)	-0.98 (1.04)	-0.68 (0.68)
Observations	2,178	8,270	3,162	453	6,489	617	3,267	8,332	4,410	617	491

Notes: This table reports sector-specific materials-emissions quadratic IV estimates of the Morishima elasticity M_{ME} between materials and emissions from equation (22). All regressions control for labor, capital, and materials prices, firm fixed effects, and a sector linear time trend. Regressions instrument for $\log(P^E)$ and $\log(P^E)^2$ using the carbon price surprise series and its square. Standard errors, from a cluster bootstrap by firm and year, are in parentheses.

B Data Construction and Measurement

This appendix describes in more detail the data cleaning steps and how we construct the empirical objects used in estimation and counterfactuals.

EUTL-Orbis Crosswalk We link installation-year European Union Transaction Log (EUTL) data to firm-year Orbis data using the European Commission/JRC EUTL-Orbis crosswalk. We aggregate total verified emissions up to the firm-year level within a firm across all of its installations.

The EUTL-Orbis linking creates a potential measurement issue. EU-ETS regulates installations based on physical activity, while Orbis assigns firms to sectors using their primary NACE code. These two can differ when a firm operates a regulated plant outside its primary business line (e.g., a cement plant operating on-site power generation). We use the Orbis sector for assignment to better reflect the economic decision-making of the firm.

Deflation and Real Variables We express firm financial variables in real terms using country-year-level price deflators from Penn World Table 11.0 (PWT). We apply this deflator to operating revenue, fixed assets, labor costs, materials expenditures, value added, and emissions prices.

Materials Prices and Quantities Orbis reports firm materials expenditures but not physical materials quantities, which we need for both estimation and simulation when constructing the productive composite. We construct materials quantities by dividing reported materials spending by the EU KLEMS intermediate input price index for the firm’s country, sector, and year.

Capital Rental Rate We construct the rental price of capital using EU KLEMS capital account data and PWT country-year returns. EU KLEMS provides investment price growth and depreciation by country, NACE section, and year. PWT provides country-year real returns and price information. We then compute the capital rental rate P_{nct}^K for sector n in country c in year t using a Jorgenson-style user cost formula:

$$P_{nct}^K = r_{ct} - (1 - \delta_{nct})g_{nct}^K + \rho.$$

$r_{ct} = 1 + \text{irr}_{ct}$ is the gross real return, which comes from the PWT net internal rate of return irr_{ct} , δ_{nct} is the depreciation rate, g_{nct}^K is one plus real investment price growth, and $\rho = 0.07$ is the risk premium. Writing $\tilde{g}_{nct}^K = g_{nct}^K - 1$ for the net growth rate, the formula is then $P_{nct}^K = \text{irr}_{ct} + \delta_{nct}(1 + \tilde{g}_{nct}^K) - \tilde{g}_{nct}^K + \rho$, which is the standard Jorgenson rental rate including depreciation. Because the PWT return is net of depreciation, the depreciation term enters exactly once.

In a small number of country-sector-year combinations, there is rapid real investment price growth, which implies a negative rental rate. In those cases we floor the rental rate at the country-year net real return irr_{ct} . EU KLEMS does not cover Norway or Switzerland, so the depreciation

and investment price components of their user costs are imputed from sector-year averages for the EU KLEMS EU20 country aggregate.

User cost formulas can imply implausibly large capital costs for some capital-intensive firm-years and push total cost above revenue, inconsistent with our model. To circumvent this issue we cap the imputed capital cost. Let $\tilde{C}_{int}^K = P_{nct}^K K_{int}$ denote the raw imputed capital cost. We set:

$$C_{int}^K = \max \left\{ 0.01 \tilde{C}_{int}^K, \min \left(\tilde{C}_{int}^K, \frac{R_{int}}{1.01} - C_{int}^L - C_{int}^M - C_{int}^E \right) \right\},$$

and then define the model rental price as $P_{int}^K = C_{int}^K / K_{int}$. The cap affects about 15.6% of firm-years and primarily affects capital-intensive sectors such as cement and power.⁴⁰

Sector Mapping We map firms to model sectors using their primary Orbis NACE code. We include 11 sectors: power, cement and glass, basic metals, refining, chemical products, pulp and paper, rubber and plastics, fabricated metals, mining and quarrying, light manufacturing, and other non-manufacturing. We choose this grouping to keep the largest directly regulated EU-ETS activities separate, such as power, cement, refining, and chemicals, while avoiding having few observations within a sector-year for estimation and counterfactuals. The NACE codes included in each sector are shown in Table A1.

Table A1 also reports coverage by sector for emissions, firms, and observations. Power accounts for 46% of total verified emissions and 1,039 unique firms. Cement and glass is the second-largest emitting sector at 15%, followed by refining (10%), basic metals (7%), and chemical products (7%).

Some EU KLEMS and Eurostat data use aggregate industry or product codes that span more than one model sector. In these cases, we allocate the aggregate flow to our 11 model sectors using observed Orbis revenue weights.

Eurostat Supply-Use Inputs We use Eurostat Supply-Use tables for two counterfactual inputs. First, intermediate use flows determine the input-output materials coefficients Ω_{mnt} . Second, final demand entries determine the sectoral final demand shares s_{nt}^F . When a country-year is missing in the data, we impute using the nearest available year for that country.

Plant-Level Electricity Data The electricity case study links large power emitters to physical generation fleets. We use the WRI Global Power Plant Database for plant ownership, fuel type, and nameplate capacity. We use ENTSO-E Transparency Platform unit-level generation data to measure annual generation and capacity factors. These plant-level data are used only for the

⁴⁰Even if we do not cap capital costs we still obtain positive sectoral average markups and thus finite demand elasticities. The only non-negligible difference when not capping capital costs is that the refining demand elasticity increases; however, this has essentially no impact on the quantitative results because its demand elasticity was already large. Our quantitative simulations can accommodate firms with costs exceeding revenues because the calibration of the model reads changes in revenues against costs, conditional on the other calibrated model parameters, as differences in TFP.

electricity sector case study and do not enter the baseline production estimates or counterfactual calibration.

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 appendix builds the justification for that approach in two layers. Section C.1 derives a general emissions-as-an-input representation result. The result does not require functional form assumptions, constant elasticities, constant returns to scale, or separability between production and abatement. Its assumptions and proof are fairly standard, but the most general model admits no clear economic interpretation for the clean technology shifters. Section C.2 supplies the missing pieces to link this result to the main text model following the by-production approach to emissions-as-an-input. It presents a separable model in which firms divert inputs from production to abatement. It then proves that the feasible set this model generates satisfies the proposition’s assumptions. In doing so, it gives clean technology B_{int} its interpretation as the productivity of abatement inputs or emissions-saving technology. The nested CES of the main text is a separate parameterization used to quantify the model.

C.1 A General Emissions-as-an-Input Representation

A central obstacle to estimation of the production side of models of emitting firms is that economists often do not observe how firms divide resources between producing output and abating emissions. Here we show that standard regularity conditions deliver that the standard and emissions-as-an-input representations share the same input requirement set for every output-emissions pair (Y_{int}, E_{int}) and the same conditional cost function. The two representations thus have the same cost-minimizing input choices. This equivalence means a supply-side representation can be estimated from observed $(K_{int}, L_{int}, M_{int}, E_{int})$ alone.

Treating emissions or environmental services as an argument of the production function is longstanding in environmental economics. Early applications include Pethig (1976) and Siebert et al. (1980), followed by Baumol and Oates (1988), Cropper and Oates (1992), Copeland and Taylor (1994) and Copeland and Taylor (2003). Pethig (2006) and Ebert and Welsch (2007) provide formal emissions-as-an-input representations under materials-balance restrictions. Murty et al. (2012), Murty (2015), and Murty and Russell (2018) develop by-production alternatives. Our approach is a special case of this latter strand of the literature. The representation argument below follows standard production theory (Shephard, 1970; McFadden, 1978). Its role here is to provide the exact object needed for our empirical application.

A firm uses capital, labor, and materials, collected in the total input vector $v_{int} = (K_{int}, L_{int}, M_{int})$, to produce output Y_{int} with emissions E_{int} as a byproduct. Firms have heterogeneous technologies. Cleaner capital, fuel switching, or process redesign may also simultaneously raise productive efficiency and lower emissions. Let Θ_{int} collect all technology shifters relevant for production and

emissions, and let $\mathcal{T}(\Theta_{int}) \subseteq \mathbb{R}_+^3 \times \mathbb{R}_+^2$ be the set of feasible tuples $(v_{int}, Y_{int}, E_{int})$. We make seven assumptions on this set. (1) $\mathcal{T}(\Theta_{int})$ contains $(0, 0, 0)$. (2) It is closed. (3) It is convex. (4) Inputs can be freely disposed of:

$$(v, Y, E) \in \mathcal{T}(\Theta), \ v' \geq v \Rightarrow (v', Y, E) \in \mathcal{T}(\Theta).$$

(5) Output is downward comprehensive:

$$(v, Y, E) \in \mathcal{T}(\Theta), \ 0 \leq Y' \leq Y \Rightarrow (v, Y', E) \in \mathcal{T}(\Theta).$$

(6) Emissions are upward comprehensive:⁴¹

$$(v, Y, E) \in \mathcal{T}(\Theta), \ E' \geq E \Rightarrow (v, Y, E') \in \mathcal{T}(\Theta).$$

(7) Finite inputs and emissions cannot generate unbounded output:

for every $(v, E) \in \mathbb{R}_+^3 \times \mathbb{R}_+$, there exists $\bar{Y}(v, E) < \infty$ such that $(v, Y, E) \in \mathcal{T}(\Theta) \Rightarrow Y \leq \bar{Y}(v, E)$.

For a vector of factor prices $p > 0$, define the associated conditional cost and the emissions-as-an-input production function by:

$$\tilde{C}(Y, E; p, \Theta) \equiv \inf\{p \cdot v : (v, Y, E) \in \mathcal{T}(\Theta)\}, \quad \tilde{G}(v, E; \Theta) \equiv \sup\{Y \geq 0 : (v, Y, E) \in \mathcal{T}(\Theta)\}.$$

From here on, we call the standard technology with emissions as a by-product the *primitive technology*. Proposition 1 establishes the conditions for there to be an equivalent emissions-as-an-input representation.

Proposition 1 *Under assumptions (1)–(7), the feasible set is the hypograph of the emissions-as-an-input production function: for every $v, E \geq 0$,*

$$(v, Y, E) \in \mathcal{T}(\Theta) \iff 0 \leq Y \leq \tilde{G}(v, E; \Theta).$$

Consequently: (i) $\tilde{G}(\cdot, \cdot; \Theta)$ is concave and upper semicontinuous on $\mathbb{R}_+^3 \times \mathbb{R}_+$ and weakly increasing in E and in v ; (ii) for every $Y, E \geq 0$ the primitive technology and the emissions-as-an-input representation have the same input requirement set, $\{v : \tilde{G}(v, E; \Theta) \geq Y\} = \{v : (v, Y, E) \in \mathcal{T}(\Theta)\}$, so their conditional costs coincide:

$$\inf_{v \geq 0} \{p \cdot v : \tilde{G}(v, E; \Theta) \geq Y\} = \tilde{C}(Y, E; p, \Theta);$$

and (iii) when (Y, E) is feasible for some input vector and prices are strictly positive, both problems attain the same nonempty set of cost-minimizing input vectors.

⁴¹Downward comprehensive means the firm could always produce less output, holding inputs and emissions fixed. Upward comprehensive means the firm could emit more, holding inputs and output fixed.

Proof: *Hypograph and regularity* Let $S(v, E; \Theta) \equiv \{Y \geq 0 : (v, Y, E) \in \mathcal{T}(\Theta)\}$. Because $(0, 0, 0) \in \mathcal{T}(\Theta)$, input free disposal and upward comprehensiveness in E imply that $0 \in S(v, E; \Theta)$ for every $v, E \geq 0$. By downward comprehensiveness, closedness, and the bounded-output condition, $S(v, E; \Theta)$ is a nonempty, bounded, closed interval of the form $[0, \tilde{G}(v, E; \Theta)]$. In particular, $\tilde{G}(v, E; \Theta)$ is finite and the supremum is attained, so the hypograph equivalence holds for every $v, E \geq 0$:

$$(v, Y, E) \in \mathcal{T}(\Theta) \iff 0 \leq Y \leq \tilde{G}(v, E; \Theta).$$

After reordering coordinates as (v, E, Y) , the set $\mathcal{T}(\Theta)$ is thus the hypograph of $\tilde{G}$. Since $\mathcal{T}(\Theta)$ is closed and convex, $\tilde{G}$ is concave and upper semicontinuous.

Monotonicity If $E' \geq E$ and $(v, Y, E) \in \mathcal{T}(\Theta)$, upward comprehensiveness gives $(v, Y, E') \in \mathcal{T}(\Theta)$. Hence $S(v, E; \Theta) \subseteq S(v, E'; \Theta)$ and therefore $\tilde{G}(v, E'; \Theta) \geq \tilde{G}(v, E; \Theta)$. The same argument with input free disposal in place of upward comprehensiveness shows $\tilde{G}$ is weakly increasing in v .

Cost equivalence. For every $Y, E \geq 0$ the hypograph equivalence gives, for each $v \geq 0$, $\tilde{G}(v, E; \Theta) \geq Y$ if and only if $(v, Y, E) \in \mathcal{T}(\Theta)$: the input requirement sets of the two representations are identical. Taking the infimum of $p \cdot v$ over this common set:

$$\inf_{v \geq 0} \{p \cdot v : \tilde{G}(v, E; \Theta) \geq Y\} = \inf_{v \geq 0} \{p \cdot v : (v, Y, E) \in \mathcal{T}(\Theta)\} = \tilde{C}(Y, E; p, \Theta).$$

Attainment Suppose $(v_0, Y, E) \in \mathcal{T}(\Theta)$ for some v_0 and every component of p is strictly positive. Restricting the minimization to $\{v : (v, Y, E) \in \mathcal{T}(\Theta), p \cdot v \leq p \cdot v_0\}$ changes neither the infimum nor the set of minimizers, since no input vector outside it can improve on v_0 . The restricted set is nonempty, closed because $\mathcal{T}(\Theta)$ is closed, and bounded because strictly positive prices cap each component of v at $p \cdot v_0$ divided by that component's price. It is then compact and $p \cdot v$ attains its minimum on it. Both problems thus attain the same infimum on the same nonempty set because the requirement sets coincide. $\square$

Proposition 1 says that $\tilde{G}$ is a well-behaved concave, upper-semicontinuous production function in which emissions behave like a productive input, and that cost-minimizing behavior under $\tilde{G}$ is identical to cost-minimizing behavior under the primitive technology. This is what permits estimating production parameters from observed inputs and emissions without observing abatement or the allocation of resources to abatement versus production.

C.2 A Separable Microfoundation

Proposition 1 does not provide an interpretation of the shifter Θ . Here we propose a separable primitive model in which firms allocate inputs between production and abatement. We prove that the feasible set it generates satisfies every assumption of the Proposition, establishing that the primitive and emissions as an input representations generate the same cost functions, and thus the same conditional input demands. This is critical to show because the conditional input demands are how we estimate the model.

In the separable model, a firm allocates capital, labor, and materials between production of output and abatement of emissions:

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

with all allocations nonnegative. Writing u for the production bundle and w for the abatement bundle, an input allocation is feasible for total inputs v_{int} when $u + w \leq v_{int}$. Let $X : \mathbb{R}_+^3 \rightarrow \mathbb{R}_+$ be a finite-valued, continuous, nondecreasing, and concave production function with $X(0) = 0$:

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

where $A_{int} > 0$ is Hicks-neutral productivity. Let $Z(\cdot; B_{int}) : \mathbb{R}_+^3 \rightarrow \mathbb{R}_+$ be a finite-valued abatement function, continuous, nondecreasing, and concave in inputs and increasing in the clean technology $B_{int} > 0$, with $Z(0; B_{int}) = 0$:

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

Better clean technology generates more abatement from the same abatement inputs, and zero inputs generate no productive composite and no abatement. Emissions per unit of the productive composite fall with abatement intensity:⁴²

$$\frac{A_{int} E_{int}}{Y_{int}} = \psi \left(\frac{A_{int} Z_{a,int}}{Y_{int}} \right). \quad (23)$$

Next, we impose several assumptions on ψ .

Assumption 8 $\psi : \mathbb{R}_+ \rightarrow (\underline{e}, \bar{e}]$ is continuous, strictly decreasing, and strictly convex, with $\psi(0) = \bar{e}$, $\lim_{a \rightarrow \infty} \psi(a) = \underline{e}$, and $0 \leq \underline{e} < \bar{e} < \infty$.

Decreasing ψ means emissions fall with abatement effort relative to production. Strict convexity means abatement runs into diminishing returns so each further cut in emissions intensity requires a larger increase in abatement intensity. Increasing marginal abatement costs in dollars follow when this is combined with concavity of X and Z , which together make the feasible set convex and the conditional cost convex in (Y, E) . The boundary conditions say that zero abatement delivers the finite maximum emissions intensity $\bar{e}$ while arbitrarily large abatement pushes intensity toward $\underline{e}$.

The primitive model generates its own feasible set. An input allocation (u, w) with $u + w \leq v$ yields output $A X(u)$ and emissions determined by equation (23). Substituting $Y = A X(u)$:

$$\mathcal{T}_{\text{prim}}(A, B) \equiv \left\{ (v, Y, E) \in \mathbb{R}_+^3 \times \mathbb{R}_+^2 : \begin{array}{l} \exists u, w \geq 0, u + w \leq v, \\ Y = A X(u), E = X(u) \psi(Z(w; B)/X(u)) \end{array} \right\},$$

⁴² A_{int}/Y_{int} on the left hand side is just $1/X(K_{y,int}, L_{y,int}, M_{y,int})$, the quantity of the productive composite used in output production.

with the convention that the emissions equality reads $E = 0$ when $X(u) = 0$ so that an allocation with no productive activity emits nothing.

The Lemma below shows that this feasible set satisfies assumptions (1)–(7) in Section C.1. Its proof needs the abatement requirement for an output-emissions pair (Y, E) , obtained by expressing equation (23) in reverse. We define this function $g_A(Y, E)$. The requirement enters the separable feasible set below as the constraint $Z(w; B) \geq g_A(Y, E)$, so properties of g_A become properties of the set. Lower semicontinuity makes the constraint survive limits and the set closed, joint convexity makes the set convex, and the fact that it is monotonic, nonincreasing in emissions and nondecreasing in output, delivers the two comprehensiveness assumptions. We first construct the requirement in emissions intensity terms, then convert it to levels and establish each property, including the boundary cases.

Under Assumption 8 the inverse ψ^{-1} is well defined on $(\underline{e}, \bar{e})$, is strictly decreasing, strictly convex, and nonnegative. It diverges to $+\infty$ as its argument falls to $\underline{e}$. Extend it to all of $\mathbb{R}$ by:

$$h(e) = \begin{cases} +\infty, & e \leq \underline{e}, \\ \psi^{-1}(e), & e \in (\underline{e}, \bar{e}), \\ 0, & e \geq \bar{e}. \end{cases}$$

The extension imposes free disposal of emissions above the zero-abatement intensity $\bar{e}$ and rules out intensities at or below $\underline{e}$. So as defined, h is nonnegative, nonincreasing, convex, and lower semicontinuous.⁴³

Inverting equation (23), the abatement required to deliver the pair (Y_{int}, E_{int}) is:

$$g_A(Y, E) \equiv \frac{Y}{A} h\left(\frac{AE}{Y}\right) \quad \text{for } Y > 0, \quad g_A(0, E) \equiv 0 \quad \text{for } E \geq 0. \quad (24)$$

Now we establish the needed properties of g_A .

Convexity for positive output For $Y > 0$, g_A is the perspective of the convex function h and is therefore jointly convex (Boyd and Vandenberghe, 2004, §3.2.6).

Monotonicity in emissions Because h is nonincreasing, g_A is nonincreasing in E .

Monotonicity in output The function g_A is also nondecreasing in Y . Rewriting the abatement requirement in terms of the emissions intensity gives, for $Y, E > 0$:

$$g_A(Y, E) = E \frac{h(e)}{e}$$

and $h(e)/e$ is nonincreasing as the product of two nonnegative nonincreasing functions of e , so higher output lowers the intensity and weakly raises the requirement. The remaining monotonicity cases involve $Y = 0$ or $E = 0$. Monotonicity follows because g_A vanishes on $Y = 0$, is nonnegative everywhere, and equals $+\infty$ whenever $E = 0 < Y$, since zero emissions with positive output would

⁴³Lower semicontinuity at $\underline{e}$ follows because $\psi^{-1}(e) \rightarrow +\infty$ as $e \downarrow \underline{e}$, and at $\bar{e}$ because $\psi^{-1}(e) \rightarrow 0$ as $e \uparrow \bar{e}$.

require an intensity at or below $\underline{e}$.

Boundary values and lower semicontinuity It remains to check the boundary $Y = 0$, where equation (24) sets g_A to zero. For $E > 0$ it is already pinned down. Once $Y \leq AE/\bar{e}$, the emissions intensity is above the zero-abatement intensity $\bar{e}$. In this case, no abatement is needed and $g_A(Y, E) = 0$. At the origin there is simply nothing to abate. At boundary points g_A equals zero, its global minimum, and a nonnegative function cannot jump below zero, so the lower semicontinuity inequality holds there automatically. On $Y > 0$, g_A is a positive continuous multiple of h evaluated at a continuous argument, so it inherits lower semicontinuity from h .

Convexity along boundary segments The final check is convexity along segments that end on the boundary. For $Y > 0$, $E, E_0 \geq 0$, and $\lambda \in (0, 1)$:

$$g_A(\lambda Y, \lambda E + (1 - \lambda)E_0) \leq g_A(\lambda Y, \lambda E) = \lambda g_A(Y, E),$$

where the inequality holds because g_A is nonincreasing in E , and the equality holds because g_A is homogeneous of degree one. Since $g_A(0, E_0) = 0$, this is the convexity inequality between the interior point (Y, E) and the boundary point $(0, E_0)$. Segments along the boundary are trivial because g_A vanishes there. Thus g_A has all the properties stated above on $\mathbb{R}_+^2$.

The separable feasible set relaxes the primitive equalities to capacity constraints:

$$\mathcal{T}_{\text{sep}}(A, B) \equiv \left\{ (v, Y, E) \in \mathbb{R}_+^3 \times \mathbb{R}_+^2 : \exists u, w \geq 0, u + w \leq v, X(u) \geq \frac{Y}{A}, Z(w; B) \geq g_A(Y, E) \right\}. \quad (25)$$

X and Z can be interpreted as capacities. A firm may leave inputs idle, produce less than its production bundle allows, or carry more abatement capacity than its emissions require.

Lemma 1 (Verification of the separable microfoundation) *Let Assumption 8 and the stated conditions on X and Z hold. (i) $\mathcal{T}_{\text{sep}}(A, B)$ satisfies assumptions (1)–(7) of Proposition 1 with $\Theta = (A, B)$. (ii) $\mathcal{T}_{\text{sep}}(A, B)$ is exactly the primitive feasible set completed by free disposal of emissions:*

$$\mathcal{T}_{\text{sep}}(A, B) = \{ (v, Y, E) \in \mathbb{R}_+^3 \times \mathbb{R}_+^2 : (v, Y, E') \in \mathcal{T}_{\text{prim}}(A, B) \text{ for some } 0 \leq E' \leq E \}.$$

In particular $\mathcal{T}_{\text{prim}}(A, B) \subseteq \mathcal{T}_{\text{sep}}(A, B)$, and the two sets coincide wherever the primitive emissions relation can hold: $Y > 0$ with $AE/Y \in (\underline{e}, \bar{e}]$, or $(Y, E) = (0, 0)$. The remaining points, with $Y = 0 < E$ or $AE/Y > \bar{e}$, sit strictly above a primitive point in emissions alone, so any positive price on emissions rules them out.

Proof: We verify assumptions (1)–(7) in turn for part (i) of the lemma, then prove part (ii).

(1) *Contains the origin* Take $u = w = 0$: $X(0) \geq 0$ and $Z(0; B) \geq 0 = g_A(0, 0)$.

(2) *Closedness* Let $(v_k, Y_k, E_k) \in \mathcal{T}_{\text{sep}}(A, B)$ converge to (v, Y, E) , with feasible input splits (u_k, w_k) . Feasibility of a projected point must be established by exhibiting a limiting feasible input split, since a projection of a closed set need not be closed. Because $0 \leq u_k, w_k \leq v_k$ and $v_k \rightarrow v$, the feasible input splits lie in a common compact box. A bounded sequence in $\mathbb{R}^n$ has a convergent

subsequence, so along a subsequence $u_k \rightarrow u^*$ and $w_k \rightarrow w^*$ with $u^* + w^* \leq v$. Continuity of X gives $X(u^*) = \lim X(u_k) \geq \lim Y_k/A = Y/A$. Continuity of Z in inputs and lower semicontinuity of g_A give:

$$Z(w^*; B) = \lim_k Z(w_k; B) \geq \liminf_k g_A(Y_k, E_k) \geq g_A(Y, E).$$

Hence $(v, Y, E) \in \mathcal{T}_{\text{sep}}(A, B)$.

(3) *Convexity* Let $(v_j, Y_j, E_j) \in \mathcal{T}_{\text{sep}}(A, B)$ with feasible input splits (u_j, w_j) for $j = 1, 2$, let $\lambda \in [0, 1]$, and let the subscript λ denote convex combinations. Then $u_\lambda + w_\lambda \leq v_\lambda$, concavity of X gives $X(u_\lambda) \geq \lambda X(u_1) + (1 - \lambda)X(u_2) \geq Y_\lambda/A$, and concavity of Z in inputs with joint convexity of g_A gives:

$$Z(w_\lambda; B) \geq \lambda Z(w_1; B) + (1 - \lambda)Z(w_2; B) \geq \lambda g_A(Y_1, E_1) + (1 - \lambda)g_A(Y_2, E_2) \geq g_A(Y_\lambda, E_\lambda).$$

(4) *Input free disposal* If $v' \geq v$, any feasible input split for (v, Y, E) is a feasible input split for (v', Y, E) .

(5) *Downward comprehensiveness in output* If $0 \leq Y' \leq Y$ then $X(u) \geq Y/A \geq Y'/A$, and $g_A(Y', E) \leq g_A(Y, E) \leq Z(w; B)$ because g_A is nondecreasing in Y .

(6) *Upward comprehensiveness in emissions* If $E' \geq E$ then $g_A(Y, E') \leq g_A(Y, E) \leq Z(w; B)$ because g_A is nonincreasing in E .

(7) *Bounded output* Any feasible input split satisfies $u \leq v$, so $Y \leq A X(u) \leq A X(v) < \infty$.

(ii) *Primitive plans satisfy the capacity constraints* We first show every primitive point is separable-feasible so the input allocation that generates it already meets both capacity constraints, with no slack. Let $(v, Y, E) \in \mathcal{T}_{\text{prim}}(A, B)$ be generated by the input allocation (u, w) . If $X(u) = 0$ then $(Y, E) = (0, 0)$, and (u, w) satisfies the constraints in equation (25) because $X(u) \geq 0$ and $Z(w; B) \geq 0 = g_A(0, 0)$. If $X(u) > 0$, output equals production capacity, so $X(u) = Y/A$ and the first constraint holds with equality. For the second, the primitive relation places the intensity $AE/Y = \psi(A Z(w; B)/Y)$ in $(\underline{e}, \bar{e}]$, the range of ψ . Inverting shows the firm's abatement capacity is exactly the requirement, $Z(w; B) = g_A(Y, E)$: on $(\underline{e}, \bar{e})$ because $h = \psi^{-1}$ there, and at $\bar{e}$ because ψ is strictly decreasing, so $\psi(A Z(w; B)/Y) = \bar{e} = \psi(0)$ forces $Z(w; B) = 0 = g_A(Y, E)$. Both capacity constraints therefore hold with equality: a primitive firm uses exactly the capacity it builds. Hence $\mathcal{T}_{\text{prim}}(A, B) \subseteq \mathcal{T}_{\text{sep}}(A, B)$, and every primitive point with positive output has intensity in $(\underline{e}, \bar{e}]$, the region on which the next step proves the converse.

Agreement where the primitive relation can hold Next we show that wherever the intensity is one that the primitive relation can generate, a separable point is itself primitive. The argument throttles slack capacity down to exact use. Let $(v, Y, E) \in \mathcal{T}_{\text{sep}}(A, B)$ with feasible input allocation (u, w) . At $(Y, E) = (0, 0)$ the allocation $(0, 0)$ generates $(v, 0, 0)$ because $X(0) = 0$. For $Y > 0$ with $AE/Y \in (\underline{e}, \bar{e}]$, the allocation may carry slack: $X(u) \geq Y/A > 0$ and $Z(w; B) \geq g_A(Y, E)$ can hold strictly. The maps $t \mapsto X(tu)$ and $s \mapsto Z(sw; B)$ are continuous on $[0, 1]$ and rise from $X(0) = 0$ and $Z(0; B) = 0$ to $X(u)$ and $Z(w; B)$, so the intermediate value theorem delivers $t^*, s^* \in [0, 1]$ with $X(t^*u) = Y/A$ and $Z(s^*w; B) = g_A(Y, E)$ exactly. The throttled allocation (t^*u, s^*w) leaves

the remaining inputs idle, produces $AX(t^*u) = Y$, is feasible because $t^*u + s^*w \leq u + w \leq v$, and satisfies equation (23) because $\psi(h(e)) = e$ for every $e \in (\underline{e}, \bar{e}]$. Hence $(v, Y, E) \in \mathcal{T}_{\text{prim}}(A, B)$: capacity a firm does not use is the same as inputs it does not allocate.

Free-disposal completion Finally we establish the separable set is the primitive set plus emissions padding, and nothing more. One direction is immediate: if $(v, Y, E') \in \mathcal{T}_{\text{prim}}(A, B)$ and $E \geq E'$, the inclusion above and upward comprehensiveness give $(v, Y, E) \in \mathcal{T}_{\text{sep}}(A, B)$. For the other direction we match each separable point to a primitive point below it in emissions alone. Let $(v, Y, E) \in \mathcal{T}_{\text{sep}}(A, B)$ with feasible input allocation (u, w) . If $Y > 0$ with $AE/Y \in (\underline{e}, \bar{e}]$, or $(Y, E) = (0, 0)$, take $E' = E$: the point is primitive by the previous step. If $Y > 0$ with $AE/Y > \bar{e}$, take $E' = \bar{e}Y/A < E$: the same allocation works for (v, Y, E') because $g_A(Y, E') = 0 \leq Z(w; B)$, so (v, Y, E') has intensity $\bar{e}$ and is primitive by the previous step. If $Y = 0 < E$, take $E' = 0$: the allocation $(0, 0)$ generates $(v, 0, 0) \in \mathcal{T}_{\text{prim}}(A, B)$. In every case $(v, Y, E') \in \mathcal{T}_{\text{prim}}(A, B)$ with $E' \leq E$, establishing the displayed identity in part (ii). $\square$

Proposition 1 applied to $\mathcal{T}_{\text{sep}}(A, B)$ delivers a concave, upper-semicontinuous production function $G(v, E; A, B)$, weakly increasing in emissions, under which cost-minimizing behavior is identical to that in the separable model. This allows estimation to proceed from observed inputs and emissions without observing the split between production and abatement.

D Model Quantification Details

This appendix specifies the full demand system, then describes how we quantify the model and solve the counterfactuals. We separate the objects recovered directly from firm-level data from those calibrated using sectoral final demand and input-output data. The baseline data are written without primes, counterfactual equilibrium objects carry primes, and hats denote relative changes between the two. We express the model in normalized form because the CES parameters are not unit-free. We express variables relative to their initial period values to anchor the production function at each firm's observed initial allocation.⁴⁴

D.1 Demand

Demand has two CES nests. Within sector n , a competitive composite producer combines the quantities Y_{int} supplied by firms $i \in \mathcal{I}_n$ into the sector quantity Y_{nt} . Let P_{int} denote firm i 's output price, ϕ_{in} its CES weight, ϵ_n the elasticity of substitution across firms, and P_{nt} the price of the sector composite. The sector quantity aggregator and unit expenditure function are:

$$Y_{nt} = \left(\sum_{i \in \mathcal{I}_n} \phi_{in}^{1/\epsilon_n} Y_{int}^{\frac{\epsilon_n-1}{\epsilon_n}} \right)^{\frac{\epsilon_n}{\epsilon_n-1}}, \quad P_{nt} = \left(\sum_{i \in \mathcal{I}_n} \phi_{in} P_{int}^{1-\epsilon_n} \right)^{\frac{1}{1-\epsilon_n}}.$$

⁴⁴See León-Ledesma et al. (2010) for an overview of why normalization is necessary for CES production.

For a given quantity of the sector composite, cost minimization gives firm demand:

$$Y_{int} = \phi_{in} Y_{nt} \left(\frac{P_{int}}{P_{nt}} \right)^{-\epsilon_n}.$$

The quantity Y_{nt} includes both household demand and materials demand from downstream firms.

Households allocate final expenditure across sector composites in a second, across-sector CES nest. Let Q_{nt}^F denote household final demand for sector n , Q_t^F the aggregate final demand quantity index, θ_{nt} the sector weight, and v the elasticity of substitution across sectors. The final demand aggregator and the final demand price index are:

$$Q_t^F = \left(\sum_{n \in \mathcal{N}} \theta_{nt}^{1/v} (Q_{nt}^F)^{\frac{v-1}{v}} \right)^{\frac{v}{v-1}}, \quad P_t^F = \left(\sum_{n \in \mathcal{N}} \theta_{nt} P_{nt}^{1-v} \right)^{\frac{1}{1-v}}. \quad (26)$$

Conditional on aggregate final demand, household demand for sector n is:

$$Q_{nt}^F = Q_t^F \theta_{nt} \left(\frac{P_{nt}}{P_t^F} \right)^{-v}.$$

Let $s_{i|n,t}^R \equiv R_{int}/S_{nt}$ denote firm i 's revenue share within sector n , where $R_{int} = P_{int}Y_{int}$ and $S_{nt} = \sum_{i \in \mathcal{I}_n} R_{int}$. Let $s_{nt}^F \equiv S_{nt}^F/S_t^F$ denote household final expenditure shares across sectors, where $S_{nt}^F = P_{nt}Q_{nt}^F$ and $S_t^F = \sum_n S_{nt}^F$.

D.2 Supply

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

$$Y_{int} = A_{int} Y_{in0} \left[(1 - \pi_{in}) X_{int}^{\frac{\sigma_n - 1}{\sigma_n}} + \pi_{in} \left(B_{int} \frac{E_{int}}{E_{in0}} \right)^{\frac{\sigma_n - 1}{\sigma_n}} \right]^{\frac{\sigma_n}{\sigma_n - 1}}, \quad (27)$$

and the inner nest is:

$$X_{int} = \left[\delta_{K,in} \left(\frac{K_{int}}{K_{in0}} \right)^{\frac{\xi_n - 1}{\xi_n}} + \delta_{L,in} \left(\frac{L_{int}}{L_{in0}} \right)^{\frac{\xi_n - 1}{\xi_n}} + \delta_{M,in} \left(\frac{M_{int}}{M_{in0}} \right)^{\frac{\xi_n - 1}{\xi_n}} \right]^{\frac{\xi_n}{\xi_n - 1}}.$$

The corresponding inner-unit cost is:

$$P_{int}^X = \left[\delta_{K,in}^{\xi_n} (K_{in0} P_{int}^K)^{1-\xi_n} + \delta_{L,in}^{\xi_n} (L_{in0} P_{int}^L)^{1-\xi_n} + \delta_{M,in}^{\xi_n} (M_{in0} P_{int}^M)^{1-\xi_n} \right]^{\frac{1}{1-\xi_n}}, \quad (28)$$

while marginal cost of output is:

$$MC_{int} = \left[(1 - \pi_{in})^{\sigma_n} (P_{int}^X)^{1-\sigma_n} + \pi_{in}^{\sigma_n} \left(\frac{E_{in0} P_t^E}{B_{int}} \right)^{1-\sigma_n} \right]^{\frac{1}{1-\sigma_n}} \frac{1}{A_{int} Y_{in0}}. \quad (29)$$

D.3 Equilibrium Pricing and Revenue

Given CES demand across firms, firms set output prices as constant markups over marginal cost:

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

Within-sector revenue satisfies:

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

where S_{nt} is total sector sales, not household final expenditure. This distinction is exactly what the input-output block introduces.

D.4 Quantification of Remaining Model Parameters

Demand weights and price normalizations We cannot separately identify firm demand shifters ϕ_{in} from firm productivity, so we normalize to equal within-sector CES weights:

$$\phi_{in} = 1/F_n \quad \text{for all } i \in \mathcal{I}_n.$$

Given observed firm revenue R_{int} , gross sector revenue S_{nt} , and aggregate gross revenue $S_t \equiv \sum_n S_{nt}$, sector and firm prices are recovered as:

$$P_{int} = P_{nt} \left[\frac{R_{int}/\phi_{in}}{S_{nt}} \right]^{\frac{1}{1-\epsilon_n}}$$

where

$$P_{nt} = \left[\frac{S_{nt}/\zeta_n}{S_t} \right]^{\frac{1}{1-v}}. \quad (30)$$

Equation (30) is a normalization used only to recover a baseline price vector from observed gross revenues; it does not pin down household expenditure shares. To recover a baseline sector price vector from observed gross sector revenue, we choose the normalization weights $\zeta_n = 1/N$. Note that these weights are not the household CES weights in equation (26). Gross sector sales and household final demand differ once materials demand is present, so ζ_n and θ_{nt} are conceptually distinct.

Household CES demand implies:

$$\theta_{nt} = s_{nt}^F \left(\frac{P_{nt}}{P_t^F} \right)^{v-1},$$

which we normalize so that $\sum_n \theta_{nt} = 1$ in each year.

Production function shares The production function share parameters are calibrated from observed cost shares:

$$\begin{aligned}\delta_{K,in} &= \frac{P_{in0}^K K_{in0}}{P_{in0}^L L_{in0} + P_{in0}^K K_{in0} + P_{in0}^M M_{in0}}, \\ \delta_{L,in} &= \frac{P_{in0}^L L_{in0}}{P_{in0}^L L_{in0} + P_{in0}^K K_{in0} + P_{in0}^M M_{in0}}, \\ \delta_{M,in} &= \frac{P_{in0}^M M_{in0}}{P_{in0}^L L_{in0} + P_{in0}^K K_{in0} + P_{in0}^M M_{in0}},\end{aligned}$$

and

$$\pi_{in} = \frac{P_0^E E_{in0}}{P_{in0}^L L_{in0} + P_{in0}^K K_{in0} + P_{in0}^M M_{in0} + P_0^E E_{in0}}.$$

Productive composite and clean technology Given the estimated inner nest elasticity ξ_n , we recover the composite unit cost P_{int}^X from equation (28) and measure the productive composite quantity as non-emissions spending deflated by that unit cost, $X_{int} = (C_{int}^K + C_{int}^L + C_{int}^M)/P_{int}^X$, so the composite spending ratio in equation (31) is observed spending. Clean technology is then obtained by inverting the outer first-order condition:

$$B_{int} = \left[\frac{P_t^E/P_0^E}{P_{int}^X/P_{in0}^X} \right] \left[\frac{1 - \pi_{in}}{\pi_{in}} \frac{(P_t^E E_{int})/(P_0^E E_{in0})}{(P_{int}^X X_{int})/(P_{in0}^X X_{in0})} \right]^{\frac{1}{\sigma_n - 1}}. \quad (31)$$

Hicks-neutral productivity Combining the markup condition with equation (29), we invert to recover the anchored Hicks-neutral scale term:

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

We do not need to separately identify A_{int} and Y_{in0} .

Baseline cost shares and materials coefficients The baseline inner nest cost shares are:

$$s_{K|X,int} = \frac{\delta_{K,in}^{\xi_n} (K_{in0} P_{int}^K)^{1 - \xi_n}}{\delta_{K,in}^{\xi_n} (K_{in0} P_{int}^K)^{1 - \xi_n} + \delta_{L,in}^{\xi_n} (L_{in0} P_{int}^L)^{1 - \xi_n} + \delta_{M,in}^{\xi_n} (M_{in0} P_{int}^M)^{1 - \xi_n}},$$

$$s_{L|X,int} = \frac{\delta_{L,in}^{\xi_n} (L_{in0} P_{int}^L)^{1-\xi_n}}{\delta_{K,in}^{\xi_n} (K_{in0} P_{int}^K)^{1-\xi_n} + \delta_{L,in}^{\xi_n} (L_{in0} P_{int}^L)^{1-\xi_n} + \delta_{M,in}^{\xi_n} (M_{in0} P_{int}^M)^{1-\xi_n}},$$

$$s_{M|X,int} = \frac{\delta_{M,in}^{\xi_n} (M_{in0} P_{int}^M)^{1-\xi_n}}{\delta_{K,in}^{\xi_n} (K_{in0} P_{int}^K)^{1-\xi_n} + \delta_{L,in}^{\xi_n} (L_{in0} P_{int}^L)^{1-\xi_n} + \delta_{M,in}^{\xi_n} (M_{in0} P_{int}^M)^{1-\xi_n}}.$$

The corresponding outer nest cost shares are:

$$s_{X,int} = \frac{(1 - \pi_{in})^{\sigma_n} (P_{int}^X)^{1-\sigma_n}}{(1 - \pi_{in})^{\sigma_n} (P_{int}^X)^{1-\sigma_n} + \pi_{in}^{\sigma_n} (E_{in0} P_t^E / B_{int})^{1-\sigma_n}}, \quad s_{E,int} = 1 - s_{X,int}.$$

Hence, firm i 's materials share in total variable cost is:

$$s_{M,int} = s_{X,int} s_{M|X,int}.$$

Sector n 's materials expenditures per dollar of sales are:

$$\alpha_{nt}^M = \frac{\epsilon_n - 1}{\epsilon_n} \sum_{i \in \mathcal{I}_n} s_{i|n,t}^R s_{M,int}. \quad (32)$$

Input-output materials bundle The input-output block maps supplier sector materials into the aggregate materials bundle used by a firm. We assign firms a Cobb-Douglas aggregator so that, for a firm in user sector n :

$$M_{int} = \varrho_{in}^{-1} \prod_{m \in \mathcal{N}} \left(\frac{M_{imnt}}{\Omega_{mnt}} \right)^{\Omega_{mnt}},$$

with supplier-by-user IO coefficients Ω_{mnt} satisfying $\sum_m \Omega_{mnt} = 1$. We calibrate these weights from Eurostat Supply-Use data. Let U_{mnt} denote intermediate use expenditure by user sector n on supplier sector m in year t . We set:

$$\Omega_{mnt} = \frac{U_{mnt}}{\sum_k U_{knt}}.$$

Ω_{mnt} determines the composition of sector n 's materials bundle across supplier sectors. The total materials spending per dollar of sector sales is governed separately by α_{nt}^M . ϱ_{in} is a firm-specific materials productivity shifter that allows firms to have different effective materials costs. The corresponding sectoral materials price index and firm-specific materials price are:

$$P_{nt}^M = \prod_{m \in \mathcal{N}} (P_{mt})^{\Omega_{mnt}}, \quad P_{int}^M = \varrho_{in} P_{nt}^M.$$

Input-output consistent baseline final demand Let $\mathbf{s}_t^F = (s_{1t}^F, \dots, s_{Nt}^F)'$ denote the sectoral final demand shares in the Eurostat data and $\mathbf{\Omega}_t$ the Eurostat input-output allocation. Given the materials coefficients in equation (32), baseline gross sector sales satisfy:

$$\mathbf{S}_t = \mathbf{S}_t^F + \mathbf{\Omega}_t \text{diag}(\boldsymbol{\alpha}_t^M) \mathbf{S}_t. \quad (33)$$

The first term on the right-hand side is household final demand for each sector's output. The second term is materials demand from downstream sectors: $\text{diag}(\alpha_t^M)S_t$ gives each user sector's total materials purchases $u_{nt} = \alpha_{nt}^M S_{nt}$, and Ω_t allocates those purchases across supplier sectors.

The sector revenues come from the firm data while Ω_t and s_t^F come from Eurostat, so no final demand vector satisfies equation (33) exactly at both the observed revenues and the Eurostat benchmarks. Aggregate final demand follows from the accounting identity $S_t^F = \sum_n S_{nt} - \sum_n u_{nt}$. We reconcile the two sources with an RAS adjustment on the flows $\Omega_{mnt}u_{nt}$ augmented with the final demand column $S_t^F s_{nt}^F$. The RAS adjustment ensures that each supplier sector's total sales sum to its observed revenue and each user sector's materials purchases sum to u_{nt} . Baseline gross sector sales then satisfy equation (33) sector by sector at the observed revenues.

Firm-specific materials prices with common proportional changes Baseline materials prices can differ across firms within the same user sector, but all firms in that sector share the same proportional change in materials prices under a counterfactual. Let baseline sectoral materials prices satisfy:

$$P_{nt}^M = \prod_m P_{mt}^{\Omega_{mnt}}, \quad P_{int}^M = \varrho_{in} P_{nt}^M,$$

and define the counterfactual materials price hat:

$$\hat{P}_{nt}^M \equiv \frac{P_{nt}^{M'}}{P_{nt}^M} = \prod_m (\hat{P}_{mt})^{\Omega_{mnt}}, \quad P_{int}^{M'} = P_{int}^M \hat{P}_{nt}^M. \quad (34)$$

D.5 Simulation of Counterfactuals

We solve the counterfactuals in relative terms. In our baseline model, capital and labor prices are fixed, so $\hat{P}_{int}^K = \hat{P}_{int}^L = 1$, while the permit price shock $\hat{P}_t^E$ is exogenous. The materials price $\hat{P}_{nt}^M$ will be endogenous. In the short run clean technology is fixed so $\hat{B}_{int} = 1$. In the long run, the change in clean technology is given by $\hat{B}_{int} = (\hat{P}_t^E)^{\hat{\gamma}_{4,n}^E}$, where $\hat{\gamma}_{4,n}^E$ is the sector's estimated five-year response from equation (21).

Given sector prices, firm inner unit cost hats are:

$$\hat{P}_{int}^X = \left[s_{K|X,int} (\hat{P}_{int}^K)^{1-\xi_n} + s_{L|X,int} (\hat{P}_{int}^L)^{1-\xi_n} + s_{M|X,int} (\hat{P}_{nt}^M)^{1-\xi_n} \right]^{\frac{1}{1-\xi_n}}, \quad (35)$$

and outer unit cost hats are:

$$\hat{c}_{int} = \left[s_{X,int} (\hat{P}_{int}^X)^{1-\sigma_n} + s_{E,int} \left(\frac{\hat{P}_t^E}{\hat{B}_{int}} \right)^{1-\sigma_n} \right]^{\frac{1}{1-\sigma_n}}. \quad (36)$$

Marginal cost divides the bundle unit cost by Hicks-neutral productivity, which is predetermined and held fixed in every counterfactual. With constant markups, $\hat{P}_{int} = \hat{c}_{int}$. Sector price hats can

be recovered by:

$$\hat{P}_{nt} = \left[\sum_{i \in \mathcal{I}_n} s_{i|n,t}^R \hat{P}_{int}^{1-\epsilon_n} \right]^{\frac{1}{1-\epsilon_n}}. \quad (37)$$

Counterfactual within-sector revenue shares are:

$$s_{i|n,t}^{R'} = \frac{s_{i|n,t}^R \hat{P}_{int}^{1-\epsilon_n}}{\sum_{j \in \mathcal{I}_n} s_{j|n,t}^R \hat{P}_{jnt}^{1-\epsilon_n}}, \quad (38)$$

and counterfactual cost shares are:

$$s'_{K|X,int} = s_{K|X,int} \left(\frac{\hat{P}_{int}^K}{\hat{P}_{int}^X} \right)^{1-\xi_n} \quad s'_{L|X,int} = s_{L|X,int} \left(\frac{\hat{P}_{int}^L}{\hat{P}_{int}^X} \right)^{1-\xi_n} \quad s'_{M|X,int} = s_{M|X,int} \left(\frac{\hat{P}_{nt}^M}{\hat{P}_{int}^X} \right)^{1-\xi_n}$$

$$s'_{X,int} = s_{X,int} \left(\frac{\hat{P}_{int}^X}{\hat{c}_{int}} \right)^{1-\sigma_n} \quad s'_{E,int} = s_{E,int} \left(\frac{\hat{P}_t^E / \hat{B}_{int}}{\hat{c}_{int}} \right)^{1-\sigma_n}.$$

This implies counterfactual materials cost shares:

$$s'_{M,int} = s'_{X,int} s'_{M|X,int}.$$

Counterfactual sectoral materials expenditures per dollar of sales are:

$$\alpha_{nt}^{M'} = \frac{\epsilon_n - 1}{\epsilon_n} \sum_{i \in \mathcal{I}_n} s_{i|n,t}^{R'} s'_{M,int}. \quad (39)$$

On the demand side, the final demand price index hat is:

$$\hat{P}_t^F = \left[\sum_n s_{nt}^F (\hat{P}_{nt})^{1-v} \right]^{\frac{1}{1-v}}, \quad (40)$$

and the counterfactual final demand expenditure shares are:

$$s_{nt}^{F'} = \frac{s_{nt}^F (\hat{P}_{nt})^{1-v}}{\sum_m s_{mt}^F (\hat{P}_{mt})^{1-v}}, \quad S_{nt}^{F'} = S_t^F s_{nt}^{F'}. \quad (41)$$

Given $\alpha_t^{M'}$ and $S_t^{F'}$, counterfactual gross sector sales are recovered from:

$$\mathbf{S}'_t = (\mathbf{I} - \mathbf{\Omega}_t \text{diag}(\alpha_t^{M'}))^{-1} \mathbf{S}_t^{F'}. \quad (42)$$

Firm revenues and output are then:

$$R'_{int} = s_{i|n,t}^{R'} S'_{nt}, \quad Y'_{int} = \frac{R'_{int}}{P'_{int}}. \quad (43)$$

The outer nest first-order condition fixes the counterfactual input ratio:

$$\frac{X'_{int}}{E'_{int}} = \frac{X_{int}}{E_{int}} \left(\frac{\hat{P}_t^E}{\hat{P}_{int}^X} \right)^{\sigma_n} \hat{B}_{int}^{1-\sigma_n},$$

and the production function in equation (27), evaluated at baseline inputs, pins down the input levels consistent with Y'_{int} . In the Leontief limit $\sigma_n = 0$ these reduce to $E'_{int} = E_{int}(Y'_{int}/Y_{int})/\hat{B}_{int}$ and $X'_{int} = X_{int}(Y'_{int}/Y_{int})$. The inner nest then splits composite spending across inputs using counterfactual cost shares:

$$K'_{int} = \frac{s'_{K|X,int} P_{int}^{X'} X'_{int}}{P_{int}^{K'}}, \quad L'_{int} = \frac{s'_{L|X,int} P_{int}^{X'} X'_{int}}{P_{int}^{L'}}, \quad M'_{int} = \frac{s'_{M|X,int} P_{int}^{X'} X'_{int}}{P_{int}^{M'}}. \quad (44)$$

D.5.1 Counterfactual Algorithm Steps

The counterfactual algorithm is therefore:

1. From the observed data, recover firm-side primitives $(\delta_{K,in}, \delta_{L,in}, \delta_{M,in}, \pi_{in}, B_{int}, A_{int} Y_{in0})$ and baseline prices $(P_{int}, P_{nt}, P_{int}^X, P_{int}^M)$, and benchmark final demand shares s_{nt}^F , materials coefficients α_{nt}^M , and the input-output consistent baseline sector sales vector $\mathbf{S}_t$.
2. Guess an initial vector of sector price hats $\hat{P}_{nt}^{(0)} = 1$.
3. Iterate on the sector price fixed point: update materials price hats from equation (34), then firm unit costs and prices from equations (35) and (36), and finally sector price hats from equation (37) until convergence.
4. Update within-sector revenue shares and materials coefficients using equations (38) and (39).
5. Update household final demand using equations (40) and (41).
6. Solve the Leontief system in equation (42) for counterfactual gross sector sales.
7. Recover firm revenues, outputs, and factor demands using equations (43) and (44).

D.5.2 Simulations Shutting Down Reallocation

The model makes it straightforward to shut down specific reallocation margins. Table 1, Appendix Table A6, and Figure 7 use the restrictions below.

To eliminate within-sector reallocation, we freeze the baseline firm shares in the sectoral quantity bundle: $Y'_{int} = a_{int} Y'_{nt}$, where $a_{int} \equiv Y_{int}/Y_{nt}$. With frozen shares, the sector price is the price of this fixed bundle:

$$P'_{nt} = \sum_{i \in \mathcal{I}_n} a_{int} P'_{int},$$

which replaces the CES index in equation (37) everywhere sector prices enter the algorithm.

To eliminate cross-sector reallocation, we freeze the baseline composition of household final demand quantities. Table 1 and Appendix Table A6 hold nominal final spending fixed, as in the algorithm above, so the frozen quantity mix is rescaled by a common factor that keeps final spending at counterfactual prices equal to its baseline level:

$$Q_{nt}^{F'} = Q_{nt}^F \frac{S_t^F}{\sum_{m \in \mathcal{N}} P_{mt}' Q_{mt}^F}.$$

This allows the aggregate scale response to the price level (through final demand quantities) to survive.

The emissions decomposition in Figure 7 imposes the same reallocation restrictions but fixes real final demand in its intermediate steps. The two exhibits differ only in whether nominal spending or real final demand is held fixed because they answer different questions. The tables compare outcomes with and without reallocation under a fixed household budget so final demand quantities can adjust. The figure splits the effects up into more granular channels using a series of intermediate steps. Its intermediate steps exclude the aggregate response for each bar to isolate a single margin. The steps with cross-sector reallocation shut down hold sectoral final demand quantities exactly at baseline, $Q_{nt}^{F'} = Q_{nt}^F$, and the steps that switch cross-sector reallocation back on hold the aggregate final demand quantity index fixed while the sectoral mix responds through CES demand, $Q_{nt}^{F'} = Q_{nt}^F (\hat{P}_{nt} / \hat{P}_t^F)^{-\nu}$. The final step relaxes fixed final demand quantities to fixed final demand expenditures. The gap between these last two steps is the aggregate demand channel, and the bars sum exactly to the same full model counterfactual reported in the tables with reallocation active.

D.6 Emissions Imputation for Non-ETS Firms

Verified emissions are observed only for firms linked to EU-ETS installations. For the non-ETS firms that enter the full dataset used to quantify leakage to unregulated firms, we impute baseline emissions intensities by matching each non-ETS firm-year to comparable ETS firm-years. For each non-ETS firm-year we restrict potential matches to ETS firm-years with positive verified emissions and a complete set of matching covariates. We then draw potential matches from the first pool that contains at least $k = 10$ eligible ETS observations in the following hierarchy of donor pools: (1) the same sector, country, and year; (2) the same sector and year, pooling across countries; and (3) the same sector, pooling across countries and years. Every sector contains at least 453 potential ETS firm-year observations, so each non-ETS firm-year with complete matching covariates and positive real revenue receives a full set of k matches. The remaining non-ETS firm-years are assigned zero imputed emissions. For each non-ETS firm-year observation, we compute the Mahalanobis distance to potential ETS firm-year observation matches using one of four sets of matching variables: (1) log capital, log labor, and log materials; (2) the capital, labor, and materials cost-to-revenue ratios; (3) log revenue; or (4) all of these variables combined. We assign the recipient the equal-weight arithmetic mean of the emissions intensities of the k nearest neighbor matches. Next, using the output prices from the model baseline, we can convert the imputed non-ETS sales-based emissions

intensity to a standard quantity emissions intensity of emissions divided by output.⁴⁵ In the counterfactuals, each unregulated firm's emissions scale with its solved output at this fixed quantity intensity, while regulated firms' emissions respond through the model's first-order conditions. As a coarser alternative to nearest-neighbor matching, we also impute using the revenue-weighted mean emissions intensity of ETS firms in the same country, sector, and year. Appendix Table A9 reports leakage rates under all five imputation strategies.

E General Equilibrium Extension

This appendix extends the model to general equilibrium, in which aggregate quantities, rather than prices, are exogenous. In this version of the model, there is fixed aggregate capital supply $\bar{K}_t$, fixed aggregate labor supply $\bar{L}_t$, and an economy-wide emissions cap $\bar{E}_t$. The rental rate r_t , permit price τ_t , and household income I_t are endogenous and adjust to clear markets. Here we use τ_t for the permit price to distinguish it from the exogenous partial equilibrium price P_t^E . Throughout, $\mathcal{N}$ denotes the set of sectors, and $\mathcal{I}_n$ denotes the set of firms in sector n . The model is identical to that in the main text except that it adds endogenous household income and the market clearing conditions for labor, capital, and emissions.

E.1 Income and Factor Market Equilibrium

The household owns the fixed labor and capital supplies, receives firm profits, and is paid the permit revenue as a lump-sum transfer T_t . Aggregate firm profits Π_t are:

$$\Pi_t = \sum_{n \in \mathcal{N}} (1 - \mu_n^{-1}) S_{nt}.$$

With w_t denoting the wage, household nominal income is:

$$I_t = w_t \bar{L}_t + r_t \bar{K}_t + \Pi_t + T_t.$$

The household spends all nominal income on the final demand bundle, so its budget constraint is:

$$P_t^F Q_t^F = I_t.$$

Let K_{int} , L_{int} , and E_{int} denote firm i 's capital, labor, and emissions. Aggregate labor and capital demands must equal their fixed supplies:

$$\sum_{n \in \mathcal{N}} \sum_{i \in \mathcal{I}_n} K_{int} = \bar{K}_t, \quad \sum_{n \in \mathcal{N}} \sum_{i \in \mathcal{I}_n} L_{int} = \bar{L}_t.$$

⁴⁵Using the original imputed sales-based intensities would be implicitly assuming that emissions are proportional to sales for non-ETS firms so that price changes alone could change emissions. Converting to a quantity-based emissions intensity generates a more intuitive proportionality to output.

The emissions cap is implemented through a competitive permit market. Feasibility and complementary slackness require:

$$\sum_{n \in \mathcal{N}} \sum_{i \in \mathcal{I}_n} E_{int} \leq \bar{E}_t, \quad \tau_t \geq 0, \quad \tau_t \left(\sum_{n \in \mathcal{N}} \sum_{i \in \mathcal{I}_n} E_{int} - \bar{E}_t \right) = 0.$$

The government returns firms' permit payments to the household:

$$T_t = \tau_t \sum_{n \in \mathcal{N}} \sum_{i \in \mathcal{I}_n} E_{int}. \quad (45)$$

When the cap binds, equation (45) simplifies to $T_t = \tau_t \bar{E}_t$.

E.2 General Equilibrium Solution Algorithm Sketch

Counterfactuals can be solved in terms of counterfactual levels of variables relative to baseline using exact hat algebra approaches. In the general equilibrium model, a counterfactual shocks the aggregate supply of emissions through the cap, and clean technology shocks enter as in the partial equilibrium counterfactuals. The baseline variables are the same as in Appendix D, except that for general equilibrium we also need sector sales per dollar of household income and the division of income across labor, capital, permit revenue, and profits. Both come from the baseline input-output system and firm cost shares.

The general equilibrium solution algorithm nests the partial equilibrium sector price fixed point inside an outer loop for the rental rate, the permit price, and household income, where the wage is not solved for since it is the numeraire. Given guesses for the rental rate and permit price, the inner step solves the same price fixed point as the partial equilibrium algorithm.

The outer step updates the aggregate variables at the converged prices. We update firm cost shares, revenue shares, and household expenditure shares, and we solve the input-output system for sector sales per dollar of income. These shares decompose household income into payments to labor, capital, permits, and profits. The income identity and the capital and permit market clearing conditions then deliver updated values of income, the rental rate, and the permit price. This can then be repeated until convergence.

F 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 4.1.1. 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 these shocks act as residual labor supply shifters.

We define labor markets at the NUTS3 region-year level. NUTS3 regions are broadly comparable to U.S. counties. 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. We omit sector subscripts n except where necessary at the end for the final regressions.

Let L_{irt} denote employment at firm i in region r at time t . Define firm employment shares within a market as:

$$\varphi_{irt} \equiv \frac{L_{irt}}{\sum_{j \in \mathcal{I}_{rt}} L_{jrt}}$$

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

$$\log L_{irt} = \beta_L \log \tilde{P}_{rt}^L + D_{rt} + \tilde{u}_{irt} \quad (46)$$

where D_{rt} is an aggregate labor demand shifter common to all firms in market (r, t) , and $\tilde{u}_{irt}$ 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 factor prices). Equation (46) should therefore be interpreted as a reduced form projection after absorbing these other components of firm labor demand into D_{rt} and $\tilde{u}_{irt}$.

Define u_{irt} as the idiosyncratic shock to labor demand growth. In first differences we have that labor demand growth is:

$$\Delta \log L_{irt} = \beta_L \Delta \log \tilde{P}_{rt}^L + \Delta D_{rt} + u_{irt}. \quad (47)$$

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

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

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

$$\Delta \log \tilde{P}_{rt}^L = \alpha \Delta \log L_{rt} + \Delta \varepsilon_{rt}. \quad (48)$$

To a first order, log changes in total employment are a share-weighted average of firm-level log changes:⁴⁶

$$\Delta \log L_{rt} = \sum_{i \in \mathcal{I}_{rt}} \varphi_{ir,t-1} \Delta \log L_{irt}. \quad (49)$$

Substituting equation (47) into equation (49) yields aggregate labor demand growth:

$$\Delta \log L_{rt} = \beta_L \Delta \log \tilde{P}_{rt}^L + \Delta D_{rt} + \sum_{i \in \mathcal{I}_{rt}} \varphi_{ir,t-1} u_{irt}. \quad (50)$$

Next, combining equations (48) and (50) and solving for wage changes gives:

$$\Delta \log \tilde{P}_{rt}^L = \frac{\alpha}{1 - \alpha\beta_L} \sum_{i \in \mathcal{I}_{rt}} \varphi_{ir,t-1} u_{irt} + \frac{\alpha}{1 - \alpha\beta_L} \Delta D_{rt} + \frac{1}{1 - \alpha\beta_L} \Delta \varepsilon_{rt}. \quad (51)$$

⁴⁶Note that this will not necessarily be true in levels.

Equation (51) shows that market wages respond to a share-weighted sum of idiosyncratic firm labor demand shocks, as well as regional demand and supply shocks. The equation shows that shocks to large firms, those with substantial employment and high $\varphi_{ir,t-1}$, can generate sizable market-wide wage movements.

Define the market-level granular shifter as the deviation of the share-weighted shock from the equal-weighted shock following Gabaix and Koijen (2024):

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

To construct the instrument, we need to recover a 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:

$$\Delta \log L_{irt} = \lambda_{3,i} + \lambda_{3,r\tilde{n}t} + \hat{u}_{irt}, \quad (52)$$

where $\lambda_{3,i}$ are firm fixed effects and $\lambda_{3,r\tilde{n}t}$ are region-SIC 3 sector-year fixed effects. The latter absorb market-level conditions, local sector-specific wage shocks, and other granular demand shocks. We define $\tilde{n}$ at the SIC 3 level, rather than using the 11 model sectors, to residualize employment and wages with a much finer sectoral partition: the GIV panel contains 458 distinct SIC 3 sectors. This ensures that the residual shocks used to construct the instrument are purged of local sector-year variation at a substantially more granular level than the model sectors.

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 reflect local market conditions but also firm-specific factors such as skill mix. To obtain a wage measure closer to the market wage $\log \tilde{P}_{rt}^L$ in our model, we extract a region-year wage component from firm wage data in two steps. Let W_{irt}^{obs} denote the observed firm wage. First, we residualize log firm wages on country-SIC 3 sector-year fixed effects:

$$\log W_{irt}^{\text{obs}} = \lambda_{4,c\tilde{n}t} + e_{ir\tilde{n}t}^w,$$

where $\lambda_{4,c\tilde{n}t}$ absorbs wage variation common to a country-SIC 3 sector-year cell, including national sector-level wage trends. Second, we average the estimated residuals across all firms observed in a region-year and use this average as the market wage:

$$\log P_{rt}^L \equiv \frac{1}{N_{rt}} \sum_{i \in \mathcal{I}_{rt}} \hat{e}_{ir\tilde{n}t}^w.$$

Averaging across the N_{rt} firms in a region-year removes idiosyncratic firm-level wage variation. The country-sector-year fixed effects in the first step remove aggregate and sectoral wage movements, leaving the regional labor-market component. When instrumenting for wages we use a leave-one-out version to avoid mechanical correlation between the instrument and the firm's own shock. In the baseline we take a conservative approach and leave out all other firms in the same sector. Appendix

A presents estimates under alternative strategies for constructing the instrument such as leaving out only the firm’s own observation in the construction of that firm’s instrument, or using firm wages instead of regional wages as the endogenous regressor.

Three assumptions are needed for identification of ξ_n . The first is that the extracted shocks $\hat{u}_{irt}$ are idiosyncratic. Conditional on the fixed effects in equation (52), they are mean-independent of the market-level demand and supply shifters ΔD_{rt} and $\Delta \varepsilon_{rt}$. The second is that the lagged shares $\varphi_{ir,t-1}$ are predetermined and uncorrelated with future idiosyncratic shocks conditional on the fixed effects, so the weights themselves introduce no endogenous variation. The third is the standard exclusion assumption. Conditional on the firm and country-year fixed effects in equation (18), other firms’ idiosyncratic shocks are orthogonal to firm i ’s factor demand shock $\varepsilon_{1,int}$ and affect its materials-labor ratio only through the regional wage.

G Input Substitution: An Electricity Market Case Study

In this appendix we investigate input substitution among the largest power-sector firms in our sample. The power sector provides a useful case study because input substitution in the form of fuel switching is directly observable in unit-level generation data. The case study therefore helps develop intuition for the low emissions-composite elasticity of substitution that we estimate, which implies that within-firm input substitution is a negligible channel for short-run emissions abatement (Section 5.1.2).

For this case study we select the 10 largest power-sector firms by their 2015 EU-ETS verified emissions.⁴⁷ For each firm, we link physical plants by matching the firm’s parent name against the ownership field of the 2018 WRI Global Power Plant Database within country.⁴⁸ We then link individual generation-unit data from the European Network of Transmission System Operators for Electricity (ENTSO-E) Transparency Platform to WRI plants by string-matching plant and unit names, and compute each unit’s annual capacity factor as ENTSO-E generation divided by WRI nameplate capacity. The merged dataset allows us to observe firm output by fuel type from 2015 to 2021 for this set of 10 large firms.

The fuzzy nature of the merge required several data-processing decisions, and the resulting dataset has some notable limitations. First, the WRI ownership field is missing for much of the capacity in some country-fuel combinations; for example, nearly all Italian hydro and French nuclear capacity lacks ownership data. We address this by computing firms’ fuel shares over thermal capacity only (coal/lignite, gas, oil, biomass), where ownership coverage is better.⁴⁹ Second, because we match plants to firms using the parent group’s name, each firm’s matched fleet is measured at the corporate-group level: it includes all plants owned by any subsidiary of the group—even when the EU-ETS account belongs to a single branch.⁵⁰ Third, the ENTSO-E unit-name match

⁴⁷2015 is the first year of our unit-level generation data.

⁴⁸This direct ownership match is needed because many EU-ETS transaction log accounts identify only the installation, not the parent company.

⁴⁹We report total matched capacity across all fuels only as context for scale.


⁵⁰In addition, we count jointly owned plants the same as solo-owned plants when calculating each firm’s total

fails for Czech units, whose identifiers are four-letter codes that do not align with any WRI name field; ČEZ therefore appears in the fleet summary in Table A15 but has no capacity-factor series in Figure A18. Fourth, the initial merge made the Spanish firm Endesa look like a gas-dominant fleet, even though its reported emissions and the public record point to a coal-heavy one. We identified a problem with the As Pontes site. WRI records only the combined-cycle gas unit added there in 2008 and omits the original $\approx 1,400$ MW coal plant, while ENTSO-E lists the coal units under the name “P.G.RODR” (for Puentes de García Rodríguez), which string matching cannot connect to “As Pontes.” We fix this by adding the coal plant to the WRI panel manually and assigning its ENTSO-E units through a manual match table. Finally, the string match assigned RWE Generation UK’s $\approx 1,600$ MW Aberthaw B coal units to a nearby 51 MW gas turbine with a similar name, so the implied coal capacity factor exceeded one. The same manual match table corrects this. Table A15 summarizes the matched generation fleet used in the case-study evidence.

Figure A18 plots capacity factors by fuel for the matched case-study fleets over the post-2015 carbon price increase, and Table A16 summarizes the change in average capacity factors after 2018. At the firms with sizable coal and gas fleets (Endesa, RWE Power, Uniper, and the UK operators), coal capacity factors fall sharply after 2018 while gas capacity factors stay flat or decline outright. The coal-dominant firms (PGE, TAURON, Vattenfall/LEAG) show much smaller declines: with little or no gas capacity, they have no within-firm fuel to switch to. The one exception is Enel, whose gas output rises roughly one-for-one with its coal decline after 2018. In 8 of the 9 firms with matched generation data, then, we see no strong evidence of within-firm input substitution, consistent with our near-zero emissions-composite elasticity of substitution. When carbon prices rise, coal output falls, but it is replaced by generation from other producers at the system level rather than by gas units within the same firm.

generation capacity.

Figure A18: Capacity Factor by Fuel for Case-Study Fleets, 2015–2021



Notes: Each firm panel shows the capacity-weighted annual capacity factor for matched ENTSO-E units, by fuel bucket. The label in each panel reports the firm's largest matched coal and gas fleet over the sample. Panels are ordered with dual-fuel fleets first. The top strip plots the contemporaneous EU-ETS annual carbon price (USD/tCO₂). The dashed vertical line marks the start of the 2018 price spike. See the text of this appendix for matching details and known omissions (notably ČEZ, whose ENTSO-E unit names do not match the WRI panel).

Table A15: Largest EU-ETS Electricity Emitters by 2015 Verified Emissions

Firm	Countries	2015 Mt CO ₂	Matched MW	Dominant fuel	Share
RWE Power	DE	86.3	18,294	Coal	74%
Vattenfall / LEAG	DE	65.3	17,222	Coal	76%
PGE GiEK	PL	61.7	16,216	Coal	94%
Enel Produzione	IT	37.5	19,358	Coal	38%
Endesa Generación	ES	24.4	15,307	Coal	52%
RWE Generation UK	GB	17.3	10,636	Gas	81%
EDF Energy UK	GB	16.9	15,024	Coal	75%
TAURON Wytwarzanie	PL	14.7	4,743	Coal	97%
ČEZ	CZ	14.5	12,515	Coal	87%
Uniper Kraftwerke (DE)	DE	13.7	8,486	Coal	74%

Notes: This table reports the 10 power firms with the largest 2015 verified CO₂ emissions in the EU-ETS. “Matched MW” is the consolidated parent-level capacity matched into the WRI Global Power Plant Database via the matching procedure described in the text of this appendix. “Dominant fuel” and “Share” report the largest single-fuel share of *thermal* (coal/lignite, gas, oil, biomass) capacity within the matched fleet. Hydro and nuclear are excluded from the share denominator because WRI owner attribution for those fuel categories is incomplete in several EU countries. See the text of this appendix for the full list of caveats, including the under-matching of Italian hydro and the manual restoration of the As Pontes coal units.

Table A16: Coal and Gas Capacity Factors Before and After the 2018 Carbon Price Increase

Firm	Coal/Lignite CF			Gas CF		
	2015–17	2018–21	Δ	2015–17	2018–21	Δ
Endesa Generación	0.50	0.16	−0.34	0.27	0.27	0.00
RWE Power	0.75	0.56	−0.20	0.12	0.13	+0.01
Uniper Kraftwerke (DE)	0.45	0.26	−0.20	0.04	0.02	−0.02
RWE Generation UK	0.33	0.04	−0.29	0.48	0.30	−0.18
EDF Energy UK	0.19	0.12	−0.07	0.46	0.30	−0.16
Enel Produzione	0.39	0.18	−0.22	0.03	0.09	+0.06
Vattenfall / LEAG	0.61	0.55	−0.07	0.33	0.38	+0.04
PGE GIEK	0.55	0.51	−0.03	0.44	0.44	0.00
TAURON Wytwarzanie	0.33	0.26	−0.07	–	–	–

Notes: This table reports mean capacity-weighted annual capacity factors for matched ENTSO-E units over 2015–2017 and 2018–2021, by firm and fuel. Firms are ordered as in Figure A18. TAURON has no matched gas units. ČEZ has no matched ENTSO-E units and is omitted.

H Marginal Abatement Cost Derivations: Firm, Sector, and Economy

This appendix derives the marginal abatement cost (MAC) results shown in Sections 3.3 and 5.4. We begin with a firm, then aggregate across firms within a sector, and then aggregate across sectors. The aggregation follows the framework of Oberfield and Raval (2021), while accounting for emissions-as-an-input, input-output pass-through of emissions price changes to the composite price, and markup heterogeneity across sectors.

Throughout the derivation we focus on two distinct objects, the MAC and the price-emissions schedule. The MAC is a cost concept. It is defined as the shadow cost of removing one more unit of emissions while holding output fixed at the corresponding level of aggregation. For a single firm, its own cost function defines its MAC directly. Since a firm's cost-minimizing choices at a given output are undistorted by the markup, the firm's market behavior tells us this cost. At higher levels of aggregation we need to introduce a least-cost planner's problem. In these cases, the cost of the marginal ton depends on how abatement is allocated across heterogeneous units. The decentralized economy tells us how the permit price and equilibrium emissions move together along the economy's price-emissions schedule. This identifies the MAC only if the market allocates abatement the way the planner would. If the market allocates abatement differently, it is because of pre-existing distortions like markups. In this case, the market permit price would not reveal the actual marginal cost of abatement. At each level of aggregation we therefore set up the planner problem to define the MAC, derive the decentralized price-emissions schedule, and establish whether the two coincide. At the firm level they coincide immediately because the firm's first-order condition sets its MAC equal to the permit price. Within a sector, the common markup across firms delivers the planner's allocation exactly. Across sectors, attainment fails because markups pass across sectors and heterogeneous markups distort consumer behavior.

H.1 Firm-Level Marginal Abatement Costs

Consider firm i in sector n at time t . We derive two target expressions: the firm's price-emissions schedule $d \log E_{int} / d \log P_t^E |_{Y_{int}, B_{int}}$ and its MAC elasticity $d \log MAC_{int} / d \log E_{int} |_{Y_{int}, B_{int}}$. The second target is defined from the firm's own emissions-constrained cost, with no planner involved. For the production technology F_{int} , this cost is:

$$\mathcal{K}_{int}(Y_{int}, E_{int}; P_{int}^X) \equiv \min_{X_{int}, E' > 0} \{ P_{int}^X X_{int} \mid Y_{int} \leq F_{int}(X_{int}, E'), E' \leq E_{int} \}.$$

The firm MAC is the cost of tightening its emissions constraint by one unit:

$$MAC_{int}(Y_{int}, E_{int}; P_{int}^X) \equiv - \frac{\partial \mathcal{K}_{int}(Y_{int}, E_{int}; P_{int}^X)}{\partial E_{int}}.$$

We now derive the components needed for both targets. The firm's outer nest technology is:

$$F_{int}(X_{int}, E_{int}) \equiv A_{int} \left[(1 - \pi_{in}) X_{int}^{(\sigma_n - 1)/\sigma_n} + \pi_{in} (B_{int} E_{int})^{(\sigma_n - 1)/\sigma_n} \right]^{\sigma_n/(\sigma_n - 1)}.$$

At fixed output, the firm chooses composite inputs and emissions to minimize expenditure:

$$\mathcal{C}_{int}(Y_{int}; P_{int}^X, P_t^E) \equiv \min_{X_{int}, E_{int} > 0} \{ P_{int}^X X_{int} + P_t^E E_{int} \mid Y_{int} \leq F_{int}(X_{int}, E_{int}) \}.$$

The first-order conditions equate the input-price ratio to the ratio of marginal products:

$$\frac{P_t^E}{P_{int}^X} = \frac{\pi_{in}}{1 - \pi_{in}} B_{int}^{(\sigma_n - 1)/\sigma_n} \left(\frac{E_{int}}{X_{int}} \right)^{-1/\sigma_n}.$$

The corresponding emissions cost share is:

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

Taking logs gives equation (5) in the main text:

$$\log \left(\frac{E_{int}}{X_{int}} \right) = \sigma_n \log \left(\frac{\pi_{in}}{1 - \pi_{in}} \right) - \sigma_n \log \left(\frac{P_t^E}{P_{int}^X} \right) + (\sigma_n - 1) \log B_{int}. \quad (53)$$

To derive the emissions-share response, write the input ratio as a function of cost shares and prices:

$$\log \left(\frac{E_{int}}{X_{int}} \right) = \log \left(\frac{s_{E,int}}{1 - s_{E,int}} \right) - \log \left(\frac{P_t^E}{P_{int}^X} \right), \quad (54)$$

and because $d \log[s_{E,int}/(1 - s_{E,int})] = ds_{E,int}/[s_{E,int}(1 - s_{E,int})]$, equations (53) and (54) imply:

$$\frac{ds_{E,int}}{d \log(P_t^E/P_{int}^X)} = -(\sigma_n - 1) s_{E,int} (1 - s_{E,int}). \quad (55)$$

Next define the composite pass-through ι_{int}^X , the pass-through of the permit price into the firm's composite input price through the input-output network:

$$\iota_{int}^X \equiv \frac{d \log P_{int}^X}{d \log P_t^E}, \quad 1 - \iota_{int}^X > 0.$$

Its complement $1 - \iota_{int}^X$ is the relative price pass-through: the share of a permit price increase that survives in the firm's relative input price P_t^E/P_{int}^X . The two pass-throughs sum to one, so what passes into the composite price does not pass into the relative price. Shephard's lemma, including the input-output induced movement in P_{int}^X , gives the cost response $d \log \mathcal{C}_{int}/d \log P_t^E = s_{E,int} + (1 - s_{E,int}) \iota_{int}^X$. At fixed output and technology, the expenditure identity $E_{int} = s_{E,int} \mathcal{C}_{int}/P_t^E$ then

gives:

$$\frac{d \log E_{int}}{d \log P_t^E} = \frac{d \log s_{E,int}}{d \log P_t^E} + \frac{d \log \mathcal{C}_{int}}{d \log P_t^E} - 1. \quad (56)$$

Equation (55) and the chain rule give $d \log s_{E,int} / d \log P_t^E = -(\sigma_n - 1)(1 - s_{E,int})(1 - \iota_{int}^X)$. Substituting the share and cost responses into equation (56) yields the first target, the firm's emissions-price elasticity:

$$\left. \frac{d \log E_{int}}{d \log P_t^E} \right|_{Y_{int}, B_{int}} = -\sigma_n(1 - s_{E,int})(1 - \iota_{int}^X). \quad (57)$$

Because F_{int} is increasing in emissions and $P_{int}^X > 0$, the emissions constraint in $\mathcal{K}_{int}$ binds and $E' = E_{int}$. The MAC is the shadow price of this constraint and tells us the additional composite-input spending needed to hold output at Y_{int} when the firm must emit one unit less.⁵¹ The constrained and priced problems are connected by:

$$\mathcal{C}_{int}(Y_{int}; P_{int}^X, P_t^E) = \min_{E_{int} > 0} \{ \mathcal{K}_{int}(Y_{int}, E_{int}; P_{int}^X) + P_t^E E_{int} \},$$

so the priced firm chooses its own emissions constraint. The interior first-order condition is $MAC_{int}(Y_{int}, E_{int}; P_{int}^X) = P_t^E$: at every decentralized emissions choice, the shadow price of the constraint equals the permit price. Inverting equation (57) then gives the second target, the elasticity of the MAC with respect to emissions:

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

The elasticity is the inverse of the elasticity of substitution with two adjustments. The first is adjusting it by the composite share of firm costs $(1 - s_{E,int})$. This is to handle how the MAC elasticity is a level concept while the elasticity of substitution is a ratio concept. Economically, a firm with high emissions cost share already has a limited ability to substitute toward the composite, which gives it a steeper MAC. The second is adjusting for the input-output network. The relative price pass-through $(1 - \iota_{int}^X)$ captures the firm's exposure to higher materials and composite costs through the input-output network. If the pass-through ι_{int}^X is larger, less of the permit price change survives in the firm's relative input price, making substitution more difficult and the MAC steeper.

⁵¹Output is held fixed because the MAC is a partial derivative at fixed Y_{int} .

H.1.1 A Closed Form Firm MAC

The elasticity above characterizes the local slope. We now derive the closed form MAC.⁵² Solving the output constraint for X_{int} and then multiplying both sides by P_{int}^X gives:

$$\mathcal{K}_{int}(Y_{int}, E_{int}; P_{int}^X) = P_{int}^X \left(\frac{(Y_{int}/A_{int})^{(\sigma_n-1)/\sigma_n} - \pi_{in}(B_{int}E_{int})^{(\sigma_n-1)/\sigma_n}}{1 - \pi_{in}} \right)^{\sigma_n/(\sigma_n-1)}. \quad (58)$$

Differentiating equation (58) with respect to $-E_{int}$ produces the closed form MAC:

$$\begin{aligned} MAC_{int}(Y_{int}, E_{int}; P_{int}^X) &= P_{int}^X \frac{\pi_{in}}{1 - \pi_{in}} B_{int}^{(\sigma_n-1)/\sigma_n} E_{int}^{-1/\sigma_n} \\ &\times \left(\frac{(Y_{int}/A_{int})^{(\sigma_n-1)/\sigma_n} - \pi_{in}(B_{int}E_{int})^{(\sigma_n-1)/\sigma_n}}{1 - \pi_{in}} \right)^{1/(\sigma_n-1)}. \end{aligned} \quad (59)$$

Define the marginal cost of production $MC_{int}(Y_{int}, E_{int}) \equiv \partial \mathcal{K}_{int}(Y_{int}, E_{int}; P_{int}^X) / \partial Y_{int}$ and the output elasticity of emissions $\eta_{E,int} \equiv \partial \log F_{int} / \partial \log E_{int}$. Differentiate equation (58) with respect to Y_{int} , and differentiate F_{int} with respect to E_{int} to get:

$$MAC_{int}(Y_{int}, E_{int}; P_{int}^X) = MC_{int}(Y_{int}, E_{int}) \eta_{E,int} \frac{Y_{int}}{E_{int}}, \quad \eta_{E,int} = \pi_{in} \left(\frac{A_{int} B_{int} E_{int}}{Y_{int}} \right)^{(\sigma_n-1)/\sigma_n}.$$

Notice that the firm first-order conditions, with λ_{int} the multiplier on the output constraint, give: $P_{int}^X = \lambda_{int} \partial F_{int} / \partial X_{int}$ and $P_t^E = \lambda_{int} \partial F_{int} / \partial E_{int}$. Constant returns gives:

$$\frac{\partial F_{int}}{\partial X_{int}} X_{int} + \frac{\partial F_{int}}{\partial E_{int}} E_{int} = Y_{int},$$

so we have that the output elasticity is the emissions cost share:

$$s_{E,int} = \frac{P_t^E E_{int}}{P_{int}^X X_{int} + P_t^E E_{int}} = \frac{(\partial F_{int} / \partial E_{int}) E_{int}}{Y_{int}} = \eta_{E,int}.$$

The MAC, at the equilibrium firm choices, can then be written as the product of the firm marginal cost, emissions cost share, and inverse emissions intensity:

$$MAC_{int} = MC_{int} s_{E,int} \frac{Y_{int}}{E_{int}}.$$

⁵²The closed forms below are valid on the admissible domain of emissions levels, where the numerator in the constrained-cost expression below is strictly positive. The domain boundary is $E_{int}^{X=0} \equiv Y_{int} / [A_{int} B_{int} \pi_{in}^{\sigma_n/(\sigma_n-1)}]$. When $\sigma_n > 1$, admissible emissions satisfy $0 < E_{int} < E_{int}^{X=0}$: at $E_{int}^{X=0}$ the firm could produce Y_{int} with no composite input, so there is no emissions-composite tradeoff and the MAC is zero. When $0 < \sigma_n < 1$, both inputs are essential, $E_{int}^{X=0}$ is the minimum emissions requirement for output Y_{int} , and admissible emissions satisfy $E_{int} > E_{int}^{X=0}$.

H.1.2 How Clean Technology Affects the MAC

The closed form also shows how clean technology shifts the fixed-output MAC schedule. Differentiate equation (59) with respect to $\log B_{int}$ at fixed (Y_{int}, E_{int}) :

$$\frac{\partial \log MAC_{int}(Y_{int}, E_{int}; P_{int}^X)}{\partial \log B_{int}} = \frac{(\sigma_n - 1)(Y_{int}/A_{int})^{(\sigma_n - 1)/\sigma_n} - \sigma_n \pi_{in}(B_{int}E_{int})^{(\sigma_n - 1)/\sigma_n}}{\sigma_n [(Y_{int}/A_{int})^{(\sigma_n - 1)/\sigma_n} - \pi_{in}(B_{int}E_{int})^{(\sigma_n - 1)/\sigma_n}]}.$$

The denominator is positive on the admissible domain of emissions levels. When $0 < \sigma_n < 1$, the empirically relevant range, the numerator is negative and cleaner technology shifts the fixed-output MAC schedule downward at all E_{int} . When $\sigma_n > 1$, the sign can reverse if $\pi_{in}(B_{int}E_{int})^{(\sigma_n - 1)/\sigma_n}$ is small relative to $(Y_{int}/A_{int})^{(\sigma_n - 1)/\sigma_n}$.

H.2 Sector-Level Marginal Abatement Costs

We next derive the sector aggregate. To define the sector MAC we pose the problem of a sector planner who allocates output and emissions across firms at least cost. We show below that the decentralized allocation attains the planner's minimum, so the observed equilibrium lies on this cost function and the MAC can be recovered from market outcomes.

H.2.1 A Sector Planner

Let $\mathcal{K}_{int}(Y_{int}, E_{int})$ be the firm emissions-constrained cost in equation (58). The sector planner's emissions-constrained cost minimization problem is:

$$\begin{aligned} \mathcal{K}_{nt}(\bar{Y}_{nt}, E_{nt}) &\equiv \min_{\{Y_{int}, E_{int}\}_i} \sum_i \mathcal{K}_{int}(Y_{int}, E_{int}) \\ \text{subject to } \mathcal{Y}_n(\mathbf{Y}_{nt}) &\geq \bar{Y}_{nt}, \quad \sum_i E_{int} \leq E_{nt}, \end{aligned}$$

where $\mathcal{Y}_n$ is the sectoral composite CES aggregator over firm outputs:

$$\mathcal{Y}_n(\mathbf{Y}_{nt}) \equiv \left(\sum_i \phi_{in}^{1/\epsilon_n} Y_{int}^{(\epsilon_n - 1)/\epsilon_n} \right)^{\epsilon_n / (\epsilon_n - 1)}.$$

Both constraints bind because marginal cost and the firm MAC are positive, so the first-order conditions give positive multipliers λ_{nt} and τ_{nt} satisfying:

$$\frac{\partial \mathcal{K}_{int}}{\partial Y_{int}} = \lambda_{nt} \frac{\partial \mathcal{Y}_n}{\partial Y_{int}}, \quad -\frac{\partial \mathcal{K}_{int}}{\partial E_{int}} = \tau_{nt} \quad \text{for every } i.$$

The second condition equalizes firm MACs within a sector. The constrained envelope theorem gives the sector MAC as the shadow value of the sector emissions constraint:

$$MAC_{nt} \equiv -\frac{\partial \mathcal{K}_{nt}}{\partial E_{nt}} = \tau_{nt}. \tag{60}$$

With the split across firms already optimal, the shadow value τ_{nt} is the cost of the sector's last ton at the input prices the sector faces, which is what gives the sector MAC the definition of a marginal cost.

H.2.2 The Decentralized Allocation

As with the firm, we aim to derive expressions for the sector elasticity of substitution and the fixed-output sector price-emissions schedule. Let $E_{nt} \equiv \sum_i E_{int}$ denote sector emissions. X_{nt} and P_{nt}^X are the sector composite quantity and price indexes, and we defer their construction until later. The sector elasticity of substitution and price-emissions schedule are:

$$\bar{\sigma}_{nt}^{\text{exact}} \equiv - \frac{d \log(E_{nt}/X_{nt})}{d \log(P_t^E/P_{nt}^X)} \bigg|_{\bar{Y}_{nt}}, \quad \frac{d \log E_{nt}}{d \log P_t^E} \bigg|_{\bar{Y}_{nt}}.$$

Return to the case where firms purchase emissions permits at the permit price. To begin the expenditure decomposition, define firm variable cost, its share of sector cost, its emissions cost share, and the sector-wide emissions cost share:

$$VC_{int} \equiv P_t^E E_{int} + P_{int}^X X_{int}, \quad w_{i|n,t} \equiv \frac{VC_{int}}{\sum_j VC_{jnt}}, \quad s_{E,int} \equiv \frac{P_t^E E_{int}}{VC_{int}}, \quad \Gamma_{nt} \equiv \sum_i w_{i|n,t} s_{E,int} \in (0, 1).$$

Let $VC_{nt} \equiv \sum_i VC_{int}$ denote total sector variable cost. The definition of Γ_{nt} gives the emissions-expenditure identity $E_{nt} = \Gamma_{nt} VC_{nt} / P_t^E$, and log differentiation gives:

$$d \log E_{nt} = \frac{d \Gamma_{nt}}{\Gamma_{nt}} + d \log VC_{nt} - d \log P_t^E.$$

The composite spending identity creates the need for the sector composite price indexes:

$$P_{nt}^X X_{nt} \equiv \sum_i P_{int}^X X_{int} = (1 - \Gamma_{nt}) VC_{nt}.$$

Firm shares of sector composite spending are:

$$\omega_{i|n,t}^X \equiv \frac{w_{i|n,t}(1 - s_{E,int})}{1 - \Gamma_{nt}} = \frac{P_{int}^X X_{int}}{\sum_j P_{jnt}^X X_{jnt}}.$$

These weights define the local Divisia price and quantity differentials:⁵³

$$d \log P_{nt}^X \equiv \sum_i \omega_{i|n,t}^X d \log P_{int}^X, \quad d \log X_{nt} \equiv \sum_i \omega_{i|n,t}^X d \log X_{int}.$$

⁵³We define the index in differentials rather than levels because the elasticity below is a local object and only log changes of the aggregates enter the derivation. The share-weighted differentials are the unique local definition where the price and quantity indexes sum to the change in actual sector composite spending, which is what lets total sector cost cancel from the quantity ratio below.

These share-weighted differentials are exactly what lets actual composite spending split locally into price and quantity. Log differentiation of the composite spending identity therefore gives:

$$d \log X_{nt} = -\frac{d\Gamma_{nt}}{1 - \Gamma_{nt}} + d \log VC_{nt} - d \log P_{nt}^X.$$

Subtracting this expression from the emissions differential cancels total sector cost from the quantity ratio. Now define the sector composite pass-through $\iota_{nt}^X \equiv d \log P_{nt}^X / d \log P_t^E$, so the relative price response introduces the sector relative price pass-through $1 - \iota_{nt}^X$. The resulting identities are:

$$d \log \frac{E_{nt}}{X_{nt}} = \frac{d\Gamma_{nt}}{\Gamma_{nt}(1 - \Gamma_{nt})} - d \log \frac{P_t^E}{P_{nt}^X}, \quad d \log \frac{P_t^E}{P_{nt}^X} = (1 - \iota_{nt}^X) d \log P_t^E. \quad (61)$$

These identities show that the elasticity requires us to derive the share response $d\Gamma_{nt}$ and the relative price pass-through $1 - \iota_{nt}^X$, while the schedule also requires the total-cost response $d \log VC_{nt}$.

The firm relative price pass-through $1 - \iota_{int}^X > 0$ measures the relative price change within a firm:

$$1 - \iota_{int}^X = \frac{d \log (P_t^E / P_{int}^X)}{d \log P_t^E}.$$

The higher the composite pass-through ι_{int}^X , the smaller the relative price change. The sector relative price pass-through $1 - \iota_{nt}^X$ in equation (61) inherits this structure. Dividing the Divisia price differential $d \log P_{nt}^X = \sum_i \omega_{i|n,t}^X d \log P_{int}^X$ by $d \log P_t^E$ gives $\iota_{nt}^X = \sum_i \omega_{i|n,t}^X \iota_{int}^X$, so the sector relative price pass-through is the composite-weighted mean of its firm counterparts:

$$1 - \iota_{nt}^X = \sum_i \omega_{i|n,t}^X (1 - \iota_{int}^X). \quad (62)$$

Define firm unit cost as $c_{int} \equiv VC_{int} / Y_{int}$. It equals MC_{int} at the firm's cost-minimizing allocation under constant returns to scale, and determines both cost pass-through and the reallocation of sector output. Shephard's lemma gives its permit price response:

$$\frac{d \log c_{int}}{d \log P_t^E} = s_{E,int} + (1 - s_{E,int}) \iota_{int}^X = 1 - (1 - s_{E,int})(1 - \iota_{int}^X).$$

Unit cost responses differ across firms through the composite share scaled by the relative price pass-through, $(1 - s_{E,int})(1 - \iota_{int}^X)$. Its cost-weighted mean R_{nt} determines the sector cost response, while its emissions share-weighted counterpart H_{nt} determines the within-firm contribution to the sector emissions-share response:

$$R_{nt} \equiv \sum_i w_{i|n,t} (1 - s_{E,int})(1 - \iota_{int}^X), \quad H_{nt} \equiv \sum_i w_{i|n,t} s_{E,int} (1 - s_{E,int})(1 - \iota_{int}^X).$$

The firm share response supplies the within-firm component of $d\Gamma_{nt}$. Equation (55) gives this

response to the common permit-price change:

$$\frac{ds_{E,int}}{d \log P_t^E} = (1 - \sigma_n) s_{E,int} (1 - s_{E,int}) (1 - \iota_{int}^X).$$

The reallocation component comes from CES demand across firms. Every firm in a sector sells to the same sectoral composite producer and faces the same demand elasticity ϵ_n , so all firms charge the common markup $\mu_n = \epsilon_n / (\epsilon_n - 1)$. The firm's price is $P_{int} = \mu_n c_{int}$, and its share of sector revenue equals its cost share $w_{i|n,t}$. CES demand at fixed sector output then satisfies $d \log Y_{int} / d \log P_t^E = \epsilon_n [(1 - s_{E,int})(1 - \iota_{int}^X) - R_{nt}]$. Differentiating $VC_{int} = c_{int} Y_{int}$ gives the change in firm cost weights and the required total-cost response:

$$\frac{dw_{i|n,t}}{d \log P_t^E} = (\epsilon_n - 1) w_{i|n,t} [(1 - s_{E,int})(1 - \iota_{int}^X) - R_{nt}], \quad \frac{d \log VC_{nt}}{d \log P_t^E} = 1 - R_{nt}.$$

Combining the within-firm share response and the reallocation of cost weights gives the required sector emissions-share response. Differentiating $\Gamma_{nt} = \sum_i w_{i|n,t} s_{E,int}$ and using $\text{Cov}_w(s_{E,int}, (1 - s_{E,int})(1 - \iota_{int}^X)) = H_{nt} - \Gamma_{nt} R_{nt}$ gives:

$$\frac{d \Gamma_{nt}}{d \log P_t^E} = (1 - \sigma_n) H_{nt} + (\epsilon_n - 1) \text{Cov}_w(s_{E,int}, (1 - s_{E,int})(1 - \iota_{int}^X)) = \epsilon_n H_{nt} - \sigma_n H_{nt} - (\epsilon_n - 1) \Gamma_{nt} R_{nt}. \quad (63)$$

Combining the share and total-cost responses with $E_{nt} = \Gamma_{nt} VC_{nt} / P_t^E$ gives the price-emissions response and defines this weight:

$$\frac{d \log E_{nt}}{d \log P_t^E} = - [\zeta_{nt} \sigma_n + (1 - \zeta_{nt}) \epsilon_n] R_{nt},$$

where $\zeta_{nt} \equiv \frac{H_{nt}}{\Gamma_{nt} R_{nt}}$. To express the substitution weight ζ_{nt} in observable sector objects, we introduce two measures of heterogeneity. The dispersion index χ_{nt} records heterogeneity in emissions cost shares, while Δ_{nt}^X records the covariance between these shares and firm relative price pass-throughs:

$$\chi_{nt} \equiv \frac{\sum_i w_{i|n,t} (s_{E,int} - \Gamma_{nt})^2}{\Gamma_{nt} (1 - \Gamma_{nt})}, \quad \Delta_{nt}^X \equiv \frac{\text{Cov}_{\omega^X}(s_{E,int}, 1 - \iota_{int}^X)}{\Gamma_{nt} (1 - \iota_{nt}^X)}. \quad (64)$$

Because $s_{E,int} \in [0, 1]$ implies $\mathbb{E}_w[s_{E,int}^2] \leq \mathbb{E}_w[s_{E,int}] = \Gamma_{nt}$, the variance $\mathbb{E}_w[s_{E,int}^2] - \Gamma_{nt}^2$ is at most $\Gamma_{nt}(1 - \Gamma_{nt})$, so $0 \leq \chi_{nt} \leq 1$. Expanding the variance and changing from cost weights to composite weights gives the identities needed to write H_{nt} and ζ_{nt} in these measures:

$$\sum_i w_{i|n,t} s_{E,int} (1 - s_{E,int}) = \Gamma_{nt} (1 - \Gamma_{nt}) (1 - \chi_{nt}), \quad \sum_i \omega_{i|n,t}^X s_{E,int} = \Gamma_{nt} (1 - \chi_{nt}).$$

We now express R_{nt} , H_{nt} , and ζ_{nt} in these observables. The change of weights $w_{i|n,t} (1 - s_{E,int}) = (1 - \Gamma_{nt}) \omega_{i|n,t}^X$ turns the cost-weighted mean R_{nt} into a composite-weighted mean, and equation

(62) evaluates it:

$$R_{nt} = (1 - \Gamma_{nt}) \sum_i \omega_{i|n,t}^X (1 - \iota_{int}^X) = (1 - \Gamma_{nt})(1 - \iota_{nt}^X).$$

The same change of weights applies to H_{nt} , and splitting the resulting composite-weighted mean into its level and covariance parts with $\sum_i \omega_{i|n,t}^X s_{E,int} = \Gamma_{nt}(1 - \chi_{nt})$ gives:

$$H_{nt} = (1 - \Gamma_{nt}) \sum_i \omega_{i|n,t}^X s_{E,int} (1 - \iota_{int}^X) = (1 - \Gamma_{nt}) [\Gamma_{nt}(1 - \chi_{nt})(1 - \iota_{nt}^X) + \text{Cov}_{\omega^X}(s_{E,int}, 1 - \iota_{int}^X)].$$

Dividing H_{nt} by $\Gamma_{nt}R_{nt}$ and using the definition of Δ_{nt}^X expresses the substitution weight:

$$\zeta_{nt} = \frac{H_{nt}}{\Gamma_{nt}R_{nt}} = 1 - \chi_{nt} + \Delta_{nt}^X.$$

Substituting the share response in equation (63) together with $H_{nt} = \zeta_{nt}\Gamma_{nt}R_{nt}$ and $R_{nt} = (1 - \Gamma_{nt})(1 - \iota_{nt}^X)$ into the quantity-ratio identity, the first part of equation (61), gives:

$$\frac{d \log(E_{nt}/X_{nt})}{d \log P_t^E} = \frac{(\epsilon_n - \sigma_n)H_{nt} - (\epsilon_n - 1)\Gamma_{nt}R_{nt}}{\Gamma_{nt}(1 - \Gamma_{nt})} - (1 - \iota_{nt}^X) = [(\epsilon_n - \sigma_n)\zeta_{nt} - \epsilon_n](1 - \iota_{nt}^X).$$

This is the response of the elasticity's quantity ratio. Dividing by the relative price response, the second part of equation (61), and applying the minus sign in the definition of $\bar{\sigma}_{nt}^{\text{exact}}$ yields the elasticity:

$$\bar{\sigma}_{nt}^{\text{exact}} = \zeta_{nt}\sigma_n + (1 - \zeta_{nt})\epsilon_n.$$

Substituting $\zeta_{nt} = 1 - \chi_{nt} + \Delta_{nt}^X$ and collecting terms gives:

$$\bar{\sigma}_{nt}^{\text{exact}} = (1 - \chi_{nt} + \Delta_{nt}^X)\sigma_n + (\chi_{nt} - \Delta_{nt}^X)\epsilon_n = (1 - \chi_{nt})\sigma_n + \chi_{nt}\epsilon_n + (\sigma_n - \epsilon_n)\Delta_{nt}^X. \quad (65)$$

As in the main text, the first term captures within-firm substitution and the second captures reallocation across firms. Δ_{nt}^X in the third term captures the covariance between emissions cost shares and firm relative price pass-throughs. Firms with higher emissions cost shares are more directly exposed to permit price shocks. Firms with higher relative price pass-throughs experience larger changes in their relative input price and substitute more. When the two positively covary, the largest relative price changes land on the firms holding the emissions, so there is more within-firm substitution and the weight on σ_n rises. Reallocation across firms instead responds to differences in unit cost responses which move with $(1 - s_{E,int})(1 - \iota_{int}^X)$. The same positive covariance evens these out. The dirty firm's large permit bill pairs with an unaffected composite input bill while the clean firm's small permit bill pairs with an exposed composite one. Unit cost responses converge across dirty and clean firms, reallocation does less, and the weight on ϵ_n falls. In practice, however, this third term is small.

Finally, holding $\bar{Y}_{nt}$ and all technologies fixed, substituting equation (65) and $R_{nt} = (1 - \Gamma_{nt})(1 -$

ι_{nt}^X) into the price-emissions response gives the price-emissions schedule:

$$\left. \frac{d \log E_{nt}}{d \log P_t^E} \right|_{\bar{Y}_{nt}} = -\bar{\sigma}_{nt}^{\text{exact}}(1 - \Gamma_{nt})(1 - \iota_{nt}^X). \quad (66)$$

H.2.3 The Decentralized Allocation Attains the Planner's Minimum

We now verify that the decentralized allocation satisfies the planner's first-order conditions. With respect to emissions, each firm's problem is equivalent to $\mathcal{C}_{int} = \min_{E_{int} > 0} \{\mathcal{K}_{int}(Y_{int}, E_{int}) + P_t^E E_{int}\}$, with interior first-order condition:

$$-\frac{\partial \mathcal{K}_{int}}{\partial E_{int}} = P_t^E \quad \text{for every } i.$$

Because every firm faces the same permit price, these are the planner's emissions conditions with multiplier $\tau_{nt} = P_t^E$.

With respect to output, CES demand for firm varieties is $Y_{int} = \phi_{in}(P_{int}/P_{nt})^{-\epsilon_n} \mathcal{Y}_n$, where P_{nt} is the CES price index. Inverting the demand curve and using $\partial \mathcal{Y}_n / \partial Y_{int} = (\phi_{in} \mathcal{Y}_n / Y_{int})^{1/\epsilon_n}$ gives:

$$P_{int} = P_{nt} \left(\frac{\phi_{in} \mathcal{Y}_n}{Y_{int}} \right)^{1/\epsilon_n} = P_{nt} \frac{\partial \mathcal{Y}_n}{\partial Y_{int}}.$$

Substituting the common-markup pricing rule $P_{int} = \mu_n c_{int}$ and $c_{int} = \partial \mathcal{K}_{int} / \partial Y_{int}$, which holds at the firm's optimal emissions under constant returns, gives:

$$\frac{\partial \mathcal{K}_{int}}{\partial Y_{int}} = \frac{P_{nt}}{\mu_n} \frac{\partial \mathcal{Y}_n}{\partial Y_{int}} \quad \text{for every } i,$$

the planner's output conditions with multiplier $\lambda_{nt} = P_{nt} / \mu_n$. The common markup in our setting is what delivers a single multiplier.

H.2.4 The Sector MAC Elasticity

Equation (60) identifies the permit price with the sector MAC, $MAC_{nt} = \tau_{nt} = P_t^E$. Inverting equation (66) gives the elasticity of the sector MAC with respect to sector emissions:

$$\left. \frac{d \log MAC_{nt}}{d \log E_{nt}} \right|_{\bar{Y}_{nt}} = -\frac{1}{\bar{\sigma}_{nt}^{\text{exact}}(1 - \Gamma_{nt})(1 - \iota_{nt}^X)}.$$

In the main text explanation we set the third term to zero, $\Delta_{nt}^X = 0$, so that equation (65) becomes the convex combination in equation (11):

$$\bar{\sigma}_{nt} = (1 - \chi_{nt})\sigma_n + \chi_{nt}\epsilon_n.$$

H.2.5 First-Order Clean-Technology Approximation of χ_{nt}

We finally connect the cost-share dispersion in equation (64) to the clean-technology dispersion recovered in Appendix D. The normalized production function of Appendix D implies that a firm's emissions cost share can be written as:

$$\frac{1}{s_{E,int}} = 1 + \left(\frac{1 - \pi_{in}}{\pi_{in}} \right)^{\sigma_n} \left(\frac{P_{int}^X}{P_t^E} \right)^{1-\sigma_n} \left(\frac{B_{int}}{E_{in0}} \right)^{1-\sigma_n}. \quad (67)$$

Here E_{in0} is firm i 's initial year emissions, the normalization point of the production function. This expression makes clear that clean technology is one source of cross-firm variation in emissions cost shares.

Fix a sector-year and let $\overline{\log B_{nt}} \equiv \sum_i w_{i|n,t} \log B_{int}$. Holding the other terms in equation (67) at their expansion values, the log-odds derivative is $ds_{E,int}/d \log B_{int} = (\sigma_n - 1)s_{E,int}(1 - s_{E,int})$. A first-order expansion around $\overline{\log B_{nt}}$, with the share at the expansion point equal to Γ_{nt} to first order, gives:

$$s_{E,int} - \Gamma_{nt} \approx \Gamma_{nt}(1 - \Gamma_{nt})(\sigma_n - 1) (\log B_{int} - \overline{\log B_{nt}}). \quad (68)$$

The weighted average of the linearization equals Γ_{nt} because the log deviations average to zero.

Squaring equation (68), averaging with the cost weights $w_{i|n,t}$, and dividing by $\Gamma_{nt}(1 - \Gamma_{nt})$ yields:

$$\chi_{nt} \approx (\sigma_n - 1)^2 \Gamma_{nt}(1 - \Gamma_{nt}) \text{Var}_i^s(\log B_{int}),$$

where $\text{Var}_i^s(\log B_{int}) \equiv \sum_i w_{i|n,t} (\log B_{int} - \overline{\log B_{nt}})^2$ is the cost-weighted variance of log clean technology across firms. The approximation omits higher-order terms and cross-firm variation in P_{int}^X , π_{in} , and E_{in0} .

H.3 Economy-Wide Marginal Abatement Costs

Finally, we aggregate across sectors. At the economy level, gross production serves both households and downstream producers. Economy-wide abatement cost is necessarily a resource cost and not simply firm expenditures. The resources used up are capital and labor, whether firms devote them to producing output or to abating emissions. We call these resources *primary resources*. Why do we need to make this distinction? Materials are not primary resources because they are produced within the economy, so their capital and labor content is already counted upstream where it originates, and permit payments are not resource costs because they are transfers. Summing gross composite spending across sectors would count the same primary factors once at the producer and again inside every downstream materials bill.

To simplify notation we work in embodied units, attributing to each sector's final demand the emissions and primary resources embodied in its full supply chain. We use the vertically integrated sector representation of the input-output economy following Pasinetti (1973). The emissions side follows the environmentally extended input-output accounting of Leontief (1970). The cost side assigns each unit of final demand the direct and indirect factor payments along its supply chain

similarly to the trade in value added literature (Johnson and Noguera, 2012). In vertically integrated units, final demand is paired with the primary resources and emissions generated along its full supply chain. The accounting literature supplies the embodied attribution in levels. As at the sector level, three objects organize the derivation: the planner MAC $MAC_t \equiv -\partial \mathcal{K}_t / \partial \bar{E}_t$, defined from an emissions-constrained planner problem; the decentralized price-emissions schedule, the equilibrium response $d \log E_t^D / d \log P_t^E$; and the economy-wide elasticity of substitution $\bar{\sigma}_t$, which converts the planner's schedule into the MAC elasticity. We first state the planner problem to define the MAC. Then we build the embodied accounts, derive the decentralized schedule, and ask whether the schedule measures the planner MAC. We will see that it does not because firms impose heterogeneous markups that distort the economy-wide decentralized allocation away from the economy-wide planner's minimum. We then derive the MAC itself by pricing the planner through an *augmented expenditure function*, so the MAC follows from duality. We also show how the gap between the schedule and the MAC is governed by measurable quantities.

H.3.1 An Economy-Wide Planner

Let z_t collect outputs, primary inputs, emissions, intermediate quantities, and final deliveries, and let $\mathcal{Z}_t(\mathbf{Q}_t^F)$ be the set of allocations that satisfy every firm and sector technology and goods-market clearing when household final demand is given by $\mathbf{Q}_t^F$. Let $R_t(z_t)$ denote the primary resource cost of labor and capital and let $E_t(z_t)$ denote total emissions. Households aggregate sectoral final demands with the across-sector CES aggregator in equation (26), $\mathcal{F}(\mathbf{Q}_t^F) \equiv (\sum_{n \in \mathcal{N}} \theta_{nt}^{1/v} (Q_{nt}^F)^{(v-1)/v})^{v/(v-1)}$.

The economy-wide MAC is defined by the emissions-constrained planner problem which delivers aggregate final demand at least primary-resource cost, subject to a cap on total emissions. The emissions-constrained resource cost is:

$$\mathcal{K}_t(\bar{Q}_t^F, \bar{E}_t) \equiv \min_{\mathbf{Q}_t^F, z_t} \{ R_t(z_t) : \mathcal{F}(\mathbf{Q}_t^F) \geq \bar{Q}_t^F, z_t \in \mathcal{Z}_t(\mathbf{Q}_t^F), E_t(z_t) \leq \bar{E}_t \}. \quad (69)$$

Before proceeding we need to make a set of standard regularity assumptions:

1. Production sets are closed, convex, and constant returns to scale.
2. Primary factor prices are strictly positive.
3. The cost share input-output system has spectral radius strictly less than one.
4. The cap can be strictly undershot: some feasible allocation delivers $\bar{Q}_t^F$ with emissions strictly below $\bar{E}_t$.

Both constraints bind at the minimum because delivering final demand requires resources, and emissions are productive. The first-order conditions then give positive multipliers, λ_t on aggregate final demand and τ_t on the cap. Writing mc_{nt}^R for the marginal primary-resource cost of delivering one more unit of final delivery n at the cap:

$$mc_{nt}^R = \lambda_t \frac{\partial \mathcal{F}}{\partial Q_{nt}^F} \quad \text{for every } n, \quad -\frac{\partial \mathcal{K}_{int}}{\partial E_{int}} = \tau_t \quad \text{for every } i. \quad (70)$$

The first condition allocates final demand so that deliveries trade at their relative marginal resource costs. The second equalizes the marginal resource cost of abatement across every firm in the economy and sets it to τ_t . The constrained envelope theorem then gives the MAC as the shadow value of the cap:

$$MAC_t \equiv -\frac{\partial \mathcal{K}_t}{\partial \bar{E}_t} = \tau_t. \quad (71)$$

As at the sector level, the shadow value τ_t is the resource cost of the economy's last ton.

H.3.2 Embodied Accounts

The rest of the derivation prices the decentralized allocation in embodied units. We call sector n 's delivery to households final good n . Each final good carries its entire vertically integrated supply chain, which allows the algebra to follow similarly to the sector level. To construct these accounts, begin with the input-output network. Firms charge sector markups μ_n and buy permits at the market price P_t^E . Let $S_{nt} = P_{nt} Y_{nt}$ denote gross sector sales, and let M_{nt} denote sector n 's materials quantity with price index P_{nt}^M . Its materials share of gross sales is $\alpha_{nt}^{M,D} = P_{nt}^M M_{nt} / S_{nt}$, and $\alpha_t^{M,D}$ collects these shares across sectors. Let Ω_{mnt} be the share of sector n 's materials spending supplied by sector m , and let Ω_t collect these shares. The cost-share network matrices record how materials purchases transmit costs and emissions across sectors:

$$\mathcal{A}_{mnt}^D \equiv \Omega_{mnt} \alpha_{nt}^{M,D}, \quad \mathcal{A}_t^D \equiv \Omega_t \text{diag}(\alpha_t^{M,D}), \quad \mathcal{L}_t^D \equiv (\mathbf{I} - \mathcal{A}_t^D)^{-1},$$

where $\mathcal{L}_t^D$ exists by the spectral radius assumption. Let c_{nt}^F be the unit cost of supplying final demand, so $VC_{nt}^F = c_{nt}^F Q_{nt}^F$ is the cost of final deliveries before the final good's markup and $VC_t^F = \sum_n VC_{nt}^F$ is the total cost of satisfying final demand. Collect the markups in the diagonal matrix $\mathbf{D}_\mu = \text{diag}(\mu)$, and let $\mathbf{S}_t^D$ and $\mathbf{VC}_t^F$ collect gross sector sales and final-delivery costs across sectors. Gross sales satisfy:

$$\mathbf{S}_t^D = \mathcal{L}_t^D \mathbf{D}_\mu \mathbf{VC}_t^F.$$

The direct emissions share of gross sales is $\alpha_{nt}^{E,D} = P_t^E E_{nt} / S_{nt}$, and $\alpha_t^{E,D}$ collects these shares across sectors. The sales-based share γ_t^S then attributes emissions to a dollar of gross sales, while $\tilde{\Gamma}_t$ gives embodied emissions per dollar of final delivery cost (before the final good's markup):

$$\gamma_t^S \equiv (\mathcal{L}_t^D)' \alpha_t^{E,D}, \quad \tilde{\Gamma}_t \equiv \mathbf{D}_\mu \gamma_t^S. \quad (72)$$

Permit payments can be written as:

$$P_t^E E_t^D = \tilde{\Gamma}_t' \mathbf{VC}_t^F. \quad (73)$$

Vertical integration provides a convenient accounting reading of this price system. The matrix $\mathcal{L}_t^D$ maps final expenditure into the gross sales generated throughout each final good's supply chain, while $(\mathcal{L}_t^D)' \alpha_t^{E,D}$ attributes direct and indirect emissions to final deliveries. This is analogous

to embodied factor-content and value-added accounting (Davis and Caldeira, 2010; Johnson and Noguera, 2012).

H.3.3 The Decentralized Allocation

As with the firm and sector levels, the two target expressions are the economy-wide elasticity of substitution and the decentralized price-emissions schedule. Let E_t^D denote total emissions in the decentralized allocation. Equation (73) prices these emissions as a share of final-delivery cost. The aggregate embodied emissions cost share is:

$$\Gamma_t \equiv \frac{P_t^E E_t^D}{VC_t^F}. \quad (74)$$

A tilde distinguishes an embodied object from its direct counterpart at the firm or sector level. The final-good shares collected in $\tilde{\Gamma}_t$ carry one because Γ_{nt} already denotes the sector's direct emissions cost share, while the aggregate share Γ_t has no direct counterpart and carries none. The embodied composite $\tilde{X}_t$ is the economy-wide counterpart of the firm's composite input X_{int} . It is the result of collapsing each final delivery's supply chain into a single vertically integrated producer, and aggregating the total bundle of capital, labor, and accumulated profits that producer uses alongside emissions. $\tilde{P}_t^X$ is its price index, so its value is the non-emissions share of sales for final demand: $\tilde{P}_t^X \tilde{X}_t \equiv (1 - \Gamma_t)VC_t^F$.

The elasticity of substitution and price-emissions schedule are:

$$\bar{\sigma}_t^{\text{exact}} \equiv - \left. \frac{d \log(E_t^D / \tilde{X}_t)}{d \log(P_t^E / \tilde{P}_t^X)} \right|_{\bar{Q}_t^F}, \quad \frac{d \log E_t^D}{d \log P_t^E}.$$

The definition of Γ_t in equation (74) gives the emissions-expenditure identity $P_t^E E_t^D = \Gamma_t VC_t^F$. Log differentiation gives:

$$d \log E_t^D = \frac{d \Gamma_t}{\Gamma_t} + d \log VC_t^F - d \log P_t^E. \quad (75)$$

Log differentiating the embodied composite's value identity splits the change in its spending into price and quantity:

$$d \log \tilde{X}_t = - \frac{d \Gamma_t}{1 - \Gamma_t} + d \log VC_t^F - d \log \tilde{P}_t^X. \quad (76)$$

Subtracting equation (76) from equation (75) cancels total final-delivery cost from the quantity ratio, and the relative price response carries the aggregate composite pass-through $\iota_t^X \equiv d \log \tilde{P}_t^X / d \log P_t^E$. The expenditure identities become:

$$d \log \frac{E_t^D}{\tilde{X}_t} = \frac{d \Gamma_t}{\Gamma_t(1 - \Gamma_t)} - d \log \frac{P_t^E}{\tilde{P}_t^X}, \quad d \log \frac{P_t^E}{\tilde{P}_t^X} = (1 - \iota_t^X) d \log P_t^E. \quad (77)$$

The first identity delivers the elasticity once the aggregate responses $d \Gamma_t$ and ι_t^X are known. The

schedule is the emissions differential in equation (75): dividing it by $d \log P_t^E$ gives $d \log E_t^D / d \log P_t^E = (d\Gamma_t / d \log P_t^E) / \Gamma_t + d \log VC_t^F / d \log P_t^E - 1$. Both expressions therefore rest on the same aggregate responses built from final-good responses. Disaggregating with equation (73), $\Gamma_t VC_t^F = \sum_n \tilde{\Gamma}_{nt} VC_{nt}^F$, expresses the schedule in those final-good responses:

$$\frac{d \log E_t^D}{d \log P_t^E} = -1 + \frac{1}{P_t^E E_t^D} \sum_n \left[VC_{nt}^F \frac{d\tilde{\Gamma}_{nt}}{d \log P_t^E} + \tilde{\Gamma}_{nt} \frac{dVC_{nt}^F}{d \log P_t^E} \right]. \quad (78)$$

The sum requires the embodied share responses $d\tilde{\Gamma}_t / d \log P_t^E$ and the household spending responses $dVC_{nt}^F / d \log P_t^E$; the final-good cost responses $d \log c_{nt}^F / d \log P_t^E$ and the residual pass-throughs $\tilde{\iota}_{nt}^X$ they define deliver ι_t^X ; and $d\Gamma_t$ is assembled from all of them. We derive each in turn.

Embodied emissions per dollar of sales are the sector's own direct share plus the embodied content of its materials purchases, which is a fixed point in embodied emissions per dollar of sales. The shares in equation (72) solve this relation $\gamma_t^S = \alpha_t^{E,D} + (\mathcal{A}_t^D)' \gamma_t^S$. Differentiating this yields:

$$d\gamma_t^S = d\alpha_t^{E,D} + (d\mathcal{A}_t^D)' \gamma_t^S + (\mathcal{A}_t^D)' d\gamma_t^S.$$

Collecting the $d\gamma_t^S$ terms and solving the linear system with $(\mathbf{I} - (\mathcal{A}_t^D)')^{-1} = (\mathcal{L}_t^D)'$, then scaling by the fixed markups $\tilde{\Gamma}_t = \mathbf{D}_\mu \gamma_t^S$, gives the embodied share response needed in equation (78):

$$\frac{d\tilde{\Gamma}_t}{d \log P_t^E} = \mathbf{D}_\mu (\mathcal{L}_t^D)' \left[\frac{d\alpha_t^{E,D}}{d \log P_t^E} + \left(\frac{d\mathcal{A}_t^D}{d \log P_t^E} \right)' \gamma_t^S \right]. \quad (79)$$

The response of the cost of final deliveries follows from Shephard's lemma. The log response of sector n 's unit cost is a cost share-weighted sum of its input price responses. The shares must be on a cost basis, and sales are marked-up variable cost, $S_{nt} = \mu_n VC_{nt}$, so the cost shares of emissions and materials are the sales shares scaled by the markup, $\mu_n \alpha_{nt}^{E,D}$ and $\mu_n \alpha_{nt}^{M,D}$.⁵⁴ The emissions price rises one for one with P_t^E . The Cobb-Douglas materials bundle price rises by the expenditure-weighted mean of supplier price responses, $\sum_m \Omega_{mnt} d \log P_{mt} / d \log P_t^E$, and with fixed markups supplier prices move one for one with supplier costs, $d \log P_{mt} = d \log c_{mt}^F$. Capital and labor prices are fixed and contribute nothing. For each sector we then have $d \log c_{nt}^F / d \log P_t^E = \mu_n \alpha_{nt}^{E,D} + \mu_n \alpha_{nt}^{M,D} \sum_m \Omega_{mnt} d \log c_{mt}^F / d \log P_t^E$. Each sector's response depends on its suppliers' responses, a fixed point that can be expressed in stacked form:

$$\frac{d \log c_t^F}{d \log P_t^E} = (\mathbf{I} - \mathbf{D}_\mu (\mathcal{A}_t^D)')^{-1} \mathbf{D}_\mu \alpha_t^{E,D}. \quad (80)$$

The cost response in equation (80) and the embodied share $\tilde{\Gamma}_{nt}$ are two different numbers built from the same accounts. Under marginal cost pricing they coincide: setting $\mathbf{D}_\mu = \mathbf{I}$ in equation (80)

⁵⁴Within a sector all firms charge the common markup μ_n , so firm revenue weights equal cost weights and the sector unit cost response is the cost-weighted mean of firm responses. Sector shares are then the corresponding cost-weighted firm means.

reproduces the fixed point that defines γ_t^S , so the permit price passes into delivery costs exactly through embodied permit payments. With markups the cost response is strictly larger, because the markup scales every link of the cost fixed point but enters the embodied accounts only at the final sale.⁵⁵ We label the gap between the two as a residual pass-through $\tilde{\iota}_{nt}^X$ implicitly defined by:

$$\frac{d \log c_{nt}^F}{d \log P_t^E} = \tilde{\Gamma}_{nt} + (1 - \tilde{\Gamma}_{nt}) \tilde{\iota}_{nt}^X. \quad (81)$$

Equation (81) defines $\tilde{\iota}_{nt}^X$ only residually. It is the part of the cost response not accounted for by embodied permit payments, expressed per unit of the non-emissions share so that the decomposition mirrors the firm-level response $d \log c_{int}/d \log P_t^E = s_{E,int} + (1 - s_{E,int}) \iota_{int}^X$. It is easy to see that then $1 - d \log c_{nt}^F/d \log P_t^E = (1 - \tilde{\Gamma}_{nt})(1 - \tilde{\iota}_{nt}^X)$, so the relative price change driving substitution takes the same form as at the firm and sector levels. Although $\tilde{\iota}_{nt}^X$ is only a residual, the pricing regime pins down its value. Under the planner's marginal cost pricing, the cost response equals the embodied share and $\tilde{\iota}_{nt}^X = 0$. With the decentralized equilibrium's markups, the cost response exceeds the embodied share and $\tilde{\iota}_{nt}^X > 0$.

Household demand determines how final spending reallocates across final goods. Households buy at consumer prices $P_{nt} = \mu_n c_{nt}^F$ and face fixed markups so that $d \log P_{nt} = d \log c_{nt}^F$. Their expenditure shares are the cost weights $w_{n|t} \equiv VC_{nt}^F/VC_t^F$ tilted by heterogeneous sectoral markups. These shares are $w_{n|t}^H \equiv \mu_n w_{n|t}/\bar{\mu}_t$ where $\bar{\mu}_t \equiv \sum_m w_{m|t} \mu_m$. Let P_t^H denote the household CES price index over consumer prices. Shephard's lemma gives us that its log response is the expenditure share-weighted mean of the price responses: $d \log P_t^H/d \log P_t^E = \sum_m w_{m|t}^H d \log c_{mt}^F/d \log P_t^E$. CES demand for final good n is $Q_{nt}^F = \theta_{nt} (P_{nt}/P_t^H)^{-v} \bar{Q}_t^F$ and its log response is $d \log Q_{nt}^F/d \log P_t^E = -v(d \log c_{nt}^F/d \log P_t^E - d \log P_t^H/d \log P_t^E)$. Spending on final good n combines the cost and quantity responses $d \log VC_{nt}^F = d \log c_{nt}^F + d \log Q_{nt}^F$. We can write the responses in level terms as:

$$\frac{dQ_{nt}^F}{d \log P_t^E} = -vQ_{nt}^F \left(\frac{d \log c_{nt}^F}{d \log P_t^E} - \frac{d \log P_t^H}{d \log P_t^E} \right), \quad \frac{dVC_{nt}^F}{d \log P_t^E} = VC_{nt}^F \left[(1 - v) \frac{d \log c_{nt}^F}{d \log P_t^E} + v \frac{d \log P_t^H}{d \log P_t^E} \right]. \quad (82)$$

The household price index response differs from the cost-weighted mean pass-through, and the gap is a covariance. Substituting $w_{n|t}^H = \mu_n w_{n|t}/\bar{\mu}_t$ into the definition of $d \log P_t^H/d \log P_t^E$ splits the price index response into two parts:

$$\frac{d \log P_t^H}{d \log P_t^E} = \frac{1}{\bar{\mu}_t} \sum_n w_{n|t} \mu_n \frac{d \log c_{nt}^F}{d \log P_t^E} = \sum_n w_{n|t} \frac{d \log c_{nt}^F}{d \log P_t^E} + \frac{1}{\bar{\mu}_t} \text{Cov}_w \left(\mu_n, \frac{d \log c_{nt}^F}{d \log P_t^E} \right). \quad (83)$$

The first term is the cost-weighted mean pass-through. This is the response households would perceive if they were charged marginal cost prices. The second term appears because households face markups and weight final goods by expenditure. When high-markup goods also have high pass-through, the household price index rises by more than the cost-weighted mean. This is how

⁵⁵For example, with a single sector buying materials from itself, $d \log c_t^F/d \log P_t^E = \mu \alpha_t^{E,D}/(1 - \mu \alpha_t^{M,D})$ while $\tilde{\Gamma}_t = \mu \alpha_t^{E,D}/(1 - \alpha_t^{M,D})$. The two coincide at $\mu = 1$, and the cost response is larger for $\mu > 1$.

markups distort household behavior and it becomes the markup correction below.

Final good n 's embodied permit payments are $\tilde{\Gamma}_{nt}$ per dollar of delivery cost. Emissions embodied in one unit of final delivery can then be written as $\tilde{e}_{nt} = \tilde{\Gamma}_{nt} c_{nt}^F / P_t^E$. Log differentiation gives the within-good response, the decline in embodied emissions per unit of final delivery:

$$-\frac{d \log \tilde{e}_{nt}}{d \log P_t^E} = 1 - \frac{d \log c_{nt}^F}{d \log P_t^E} - \frac{d \log \tilde{\Gamma}_{nt}}{d \log P_t^E}. \quad (84)$$

Equation (73) and equation (74) imply $\Gamma_t = \sum_n w_{n|t} \tilde{\Gamma}_{nt}$, so the aggregate embodied share is the cost-weighted mean of final good embodied shares. We now construct the aggregate residual pass-through $1 - \iota_t^X$. Subtracting equation (81) from one factors each final good's relative price decline:

$$1 - \frac{d \log c_{nt}^F}{d \log P_t^E} = 1 - \tilde{\Gamma}_{nt} - (1 - \tilde{\Gamma}_{nt}) \tilde{\iota}_{nt}^X = (1 - \tilde{\Gamma}_{nt})(1 - \tilde{\iota}_{nt}^X). \quad (85)$$

Let $\bar{p}_t \equiv \sum_n w_{n|t} d \log c_{nt}^F / d \log P_t^E$ denote the cost-weighted mean pass-through, so the cost-weighted mean of these declines is $1 - \bar{p}_t$. Applying equation (81) at the aggregate level, $\bar{p}_t = \Gamma_t + (1 - \Gamma_t) \iota_t^X$, defines ι_t^X . In product form:

$$(1 - \Gamma_t)(1 - \iota_t^X) \equiv \sum_n w_{n|t} (1 - \tilde{\Gamma}_{nt})(1 - \tilde{\iota}_{nt}^X). \quad (86)$$

The right side is the cost-weighted mean decline in final-good relative prices, the analogue of R_{nt} in the sector derivation.

Next define the emissions share-weighted counterpart of the same declines, the analogue of H_{nt} , and the emissions share-weighted mean of the within-good responses in equation (84):

$$H_t \equiv \sum_n w_{n|t} \tilde{\Gamma}_{nt} (1 - \tilde{\Gamma}_{nt})(1 - \tilde{\iota}_{nt}^X), \quad H_{\sigma,t} \equiv \sum_n w_{n|t} \tilde{\Gamma}_{nt} \left(-\frac{d \log \tilde{e}_{nt}}{d \log P_t^E} \right).$$

$H_{\sigma,t}$ has no sector analogue. At the sector level every firm's within response was its relative price decline scaled by the common technology elasticity σ_n , so the within contribution to the sector emissions response was simply $\sigma_n H_{nt}$. Economy-wide, the final goods share no common within elasticity, so the weighted response $H_{\sigma,t}$ replaces $\sigma_n H_{nt}$. Their ratio defines the average within-good elasticity, and H_t plays its sector role through ζ_t :

$$\bar{\sigma}_t^{N,D} \equiv \frac{H_{\sigma,t}}{H_t}, \quad \zeta_t \equiv \frac{H_t}{\Gamma_t(1 - \Gamma_t)(1 - \iota_t^X)}.$$

As at the sector level, share dispersion and the covariance between embodied shares and residual pass-through determine ζ_t . With composite-cost weights $\omega_{n|t}^X \equiv w_{n|t}(1 - \tilde{\Gamma}_{nt})/(1 - \Gamma_t)$, define the dispersion index χ_t and pass-through correction Δ_t^X :

$$\chi_t \equiv \frac{\text{Var}_w(\tilde{\Gamma}_{nt})}{\Gamma_t(1 - \Gamma_t)}, \quad \Delta_t^X \equiv \frac{\text{Cov}_{\omega^X}(\tilde{\Gamma}_{nt}, 1 - \tilde{\iota}_{nt}^X)}{\Gamma_t(1 - \iota_t^X)}.$$

These let us express ζ_t as:

$$\zeta_t = 1 - \chi_t + \Delta_t^X.$$

The within-good elasticity $\bar{\sigma}_t^{N,D}$ is computed from the share responses in equation (79) and so includes the upstream recomposition of the input-output network.

Now differentiate $\Gamma_t = \sum_n w_{n|t} \tilde{\Gamma}_{nt}$. The weights respond only to relative cost changes. Each weight is the ratio $w_{n|t} = VC_{nt}^F / VC_t^F$. The spending responses in equation (82) give $d \log VC_{nt}^F / d \log P_t^E = (1-v) d \log c_{nt}^F / d \log P_t^E + v d \log P_t^H / d \log P_t^E$, and their cost-weighted sum gives $d \log VC_t^F / d \log P_t^E = (1-v) \bar{p}_t + v d \log P_t^H / d \log P_t^E$. Subtracting the two, the common household index response cancels from the ratio:

$$\frac{d \log w_{n|t}}{d \log P_t^E} = \frac{d \log VC_{nt}^F}{d \log P_t^E} - \frac{d \log VC_t^F}{d \log P_t^E} = (1-v) \left(\frac{d \log c_{nt}^F}{d \log P_t^E} - \bar{p}_t \right).$$

Only a final good's cost response relative to the mean reallocates spending, so:

$$\frac{d \Gamma_t}{d \log P_t^E} = \sum_n w_{n|t} \tilde{\Gamma}_{nt} \frac{d \log \tilde{\Gamma}_{nt}}{d \log P_t^E} + (1-v) \text{Cov}_w \left(\tilde{\Gamma}_{nt}, \frac{d \log c_{nt}^F}{d \log P_t^E} \right). \quad (87)$$

Two identities express these pieces in the objects just defined. Equation (85) rewrites H_t as $\sum_n w_{n|t} \tilde{\Gamma}_{nt} (1 - d \log c_{nt}^F / d \log P_t^E)$, and equation (84) rewrites $H_{\sigma,t}$ as the same sum net of the share responses; their difference gives the first identity below. Expanding the covariance as $\sum_n w_{n|t} \tilde{\Gamma}_{nt} d \log c_{nt}^F / d \log P_t^E - \Gamma_t \bar{p}_t$, the first term equals $\Gamma_t - H_t$ by the same rewriting, and substituting $\bar{p}_t = 1 - (1 - \Gamma_t)(1 - \iota_t^X)$ from equation (86) gives the second:

$$\sum_n w_{n|t} \tilde{\Gamma}_{nt} \frac{d \log \tilde{\Gamma}_{nt}}{d \log P_t^E} = H_t - H_{\sigma,t}, \quad \text{Cov}_w \left(\tilde{\Gamma}_{nt}, \frac{d \log c_{nt}^F}{d \log P_t^E} \right) = \Gamma_t (1 - \Gamma_t) (1 - \iota_t^X) - H_t. \quad (88)$$

Substituting the spending responses in equation (82) into equation (78) and collecting terms with equations (88) and (83) gives the emissions response:

$$\frac{d \log E_t^D}{d \log P_t^E} = - \left[\zeta_t \bar{\sigma}_t^{N,D} + (1 - \zeta_t) v \right] (1 - \Gamma_t) (1 - \iota_t^X) + v \frac{\text{Cov}_w \left(\mu_n, \frac{d \log c_{nt}^F}{d \log P_t^E} \right)}{\bar{\mu}_t}. \quad (89)$$

The bracketed terms are the analogue of $\zeta_{nt} \sigma_n + (1 - \zeta_{nt}) \epsilon_n$ at the sector level. The covariance is the economy-wide household-weight wedge in equation (83).

Substituting the aggregate share response in equation (87) and the covariance identity in equation (88) into the quantity-ratio identity, the first part of equation (77), and dividing by the relative price identity in its second part, gives the exact elasticity decomposition:

$$\bar{\sigma}_t^{\text{exact}} = \zeta_t \bar{\sigma}_t^{N,D} + (1 - \zeta_t) v. \quad (90)$$

This is the analogue of the sector result $\bar{\sigma}_{nt}^{\text{exact}} = \zeta_{nt} \sigma_n + (1 - \zeta_{nt}) \epsilon_n$.

Finally, normalize the household-weight covariance by the aggregate relative price pass-through to define the markup correction:

$$\Delta_t^\mu \equiv \frac{\text{Cov}_w\left(\mu_n, \frac{d \log c_{nt}^F}{d \log P_t^E}\right)}{\left(\sum_n w_{n|t} \mu_n\right) (1 - \Gamma_t)(1 - \iota_t^X)}. \quad (91)$$

Substituting equations (90) and (91) into equation (89) gives the exact decentralized schedule:

$$\frac{d \log E_t^D}{d \log P_t^E} = -(\bar{\sigma}_t^{\text{exact}} - v \Delta_t^\mu) (1 - \Gamma_t)(1 - \iota_t^X). \quad (92)$$

H.3.4 The Decentralized Allocation Does Not Attain the Planner's Minimum

Equation (92) is the inverse slope of the decentralized permit-price schedule. We now show that it is not the slope of the planner MAC because the decentralized allocation does not attain the minimum in equation (69). At the planner's minimum, the first-order conditions in equation (70) require (1) every firm's marginal abatement to cost primary resources with value τ_t , and (2) household demand to trade final goods at their relative marginal resource costs mc_{nt}^R .

Consider the emissions condition in equation (70) first. Each firm equalizes its private MAC to P_t^E at the market prices it faces. A dollar of marginal abatement in the decentralized equilibrium buys inputs at marked-up prices. These markups are a transfer of profit along the supply chain rather than a payment for resources or for permits valued at τ_t . The profit content of that dollar differs across firms because embodied markup content differs across supply chains. Equalizing private MACs at P_t^E thus fails the emissions condition in two ways: (1) it does not equalize marginal resource costs across firms, and (2) it does not set their level to τ_t . The level failure (2) remains even with common markups while the cross-firm failure (1) requires heterogeneity.

Markups also make final good costs rise by more than their embodied permit payments. The residual pass-throughs in equation (81) capture this. Under marginal cost pricing every $\tilde{\iota}_{nt}^X$ is zero and delivery costs respond only through embodied permit payments. At the decentralized equilibrium allocation, $\tilde{\iota}_{nt}^X$ is positive because marked-up input bills carry profit (the level failure (2)), and the residual pass-throughs differ across final goods because supply chains accumulate different markups (the cross-firm failure (1)).

Finally consider the final demand condition. In a decentralized equilibrium, households equate marginal rates of substitution to consumer prices $\mu_n c_{nt}^F$. These are proportional to the marginal resource costs only if the vector of markups in all sectors leads to common changes in consumer prices across final goods. With heterogeneous markups this will not happen. The markup correction Δ_t^μ in equation (91) is the measured violation. The gap between the schedule and the planner MAC can be decomposed into $\tilde{\iota}_{nt}^X$ and Δ_t^μ . The level of $\tilde{\iota}_{nt}^X$ scales the schedule by $1 - \iota_t^X$, and the dispersion in $\tilde{\iota}_{nt}^X$ separates ζ_t from $1 - \chi_t$.

The failure is specific to the economy-wide problem. A markup distorts the level of output rather than the input mix that delivers it. Shephard's lemma still applies at the firm's own choices,

and a markup common to every firm in a sector cancels from the relative prices inside the sector problem. Marked-up input prices do accumulate inside P_{nt}^M . The difference at the firm and sector levels is that they are genuine costs to those buyers, so the firm and sector envelope arguments go through at the prices actually faced. The economy-wide planner prices primary resources, and markups differ across the choices *inside* its problem.

This failure is an instance of a general finding for aggregation in distorted economies. Hulten (1978) shows that at an efficient allocation, induced reallocation has no first-order effect so the aggregate impact of sectoral shocks is summarized by sales weights. With cross-sector markup wedges, induced reallocations can instead matter at first order (Baqae and Farhi, 2020). The markup correction Δ_t^μ here is the model-specific analogue of this effect. It is a covariance between markups and permit price pass-through that separates the price-emissions schedule from the planner's MAC.

H.3.5 The Economy-Wide MAC Elasticity

Because the decentralized allocation does not attain the planner's minimum, we compute the MAC elasticity $d \log MAC_t / d \log E_t|_{\bar{Q}_t^F}$ from the planner problem directly. To derive the MAC we hand the planner a price instead of a cap. The augmented expenditure function is the least combined primary-resource and permit cost of delivering aggregate final demand when every ton of emissions costs τ_t :

$$G_t(\bar{Q}_t^F, \tau_t) \equiv \min_{\mathbf{Q}_t^F, z_t} \{R_t(z_t) + \tau_t E_t(z_t) : \mathcal{F}(\mathbf{Q}_t^F) \geq \bar{Q}_t^F, z_t \in \mathcal{Z}_t(\mathbf{Q}_t^F)\}.$$

The priced problem is the tool that delivers the derivative of $\mathcal{K}_t$ in equation (71).

The envelope theorem gives us that the derivative of G_t with respect to the emissions price is the minimizer's emissions:

$$G_{\tau,t}(\bar{Q}_t^F, \tau_t) = E_t^*(\bar{Q}_t^F, \tau_t). \quad (93)$$

Now return to the capped problem. Under a cap, the planner selects the same allocation as when facing a price: the emissions the priced planner chooses at τ_t define a cap, $\bar{E}_t = G_{\tau,t}(\bar{Q}_t^F, \tau_t)$, and the shadow price of that cap is τ_t . By equation (71), the permit price the planner would post to induce the cap is exactly the resource cost of the last ton. Equation (93) is therefore the planner's price-emissions schedule. Wherever the schedule is strictly decreasing, the MAC curve is its local inverse, so the MAC slope is the reciprocal of the schedule's elasticity. We now derive this schedule.

The input-output network Every aggregate from here on is evaluated at the planner's allocation, where goods are priced at marginal cost and emissions at τ_t . We reuse the symbols of the decentralized accounts with that understanding. Here $VC_{nt}^F = P_{nt}Q_{nt}^F$ is the value of final deliveries to households, the embodied variable cost of delivery n at marginal cost prices, and the direct emissions and materials shares are:

$$\alpha_{nt}^E \equiv \frac{\tau_t E_{nt}}{S_{nt}}, \quad \alpha_{nt}^M \equiv \frac{P_{nt}^M M_{nt}}{S_{nt}}.$$

Define the marginal cost counterparts of the network matrices:

$$\mathcal{A}_{mnt} \equiv \Omega_{mnt} \alpha_{nt}^M, \quad \mathcal{A}_t \equiv \Omega_t \text{diag}(\alpha_t^M), \quad \mathcal{L}_t \equiv (\mathbf{I} - \mathcal{A}_t)^{-1}.$$

When the spectral radius of $\mathcal{A}_t$ is strictly less than one, goods market clearing is:

$$\mathbf{S}_t = \mathbf{V} \mathbf{C}_t^F + \mathcal{A}_t \mathbf{S}_t, \quad \mathbf{S}_t = \mathcal{L}_t \mathbf{V} \mathbf{C}_t^F. \quad (94)$$

Notice that total emissions satisfy:

$$\tau_t E_t = (\alpha_t^E)' \mathbf{S}_t = (\alpha_t^E)' \mathcal{L}_t \mathbf{V} \mathbf{C}_t^F. \quad (95)$$

The pass-through of the permit price into sector n 's price is the log derivative $d \log P_{nt} / d \log \tau_t$. Shephard's lemma gives that the log response of a sector's unit cost to the permit price is a cost share-weighted sum of its input price responses. Emissions have cost share α_{nt}^E and their price rises one for one with τ_t . Materials have cost share α_{nt}^M and their Cobb-Douglas bundle price rises by the expenditure-weighted mean of supplier pass-throughs: $\sum_m \Omega_{mnt} d \log P_{mt} / d \log \tau_t$. Capital and labor prices are fixed and so contribute nothing. Sector by sector we then have that the pass-through can be written as: $d \log P_{nt} / d \log \tau_t = \alpha_{nt}^E + \alpha_{nt}^M \sum_m \Omega_{mnt} d \log P_{mt} / d \log \tau_t$. Each sector's pass-through depends on its suppliers' pass-throughs, a fixed point that the Leontief inverse solves by summing the whole chain:

$$\frac{d \log \mathbf{P}_t}{d \log \tau_t} = \alpha_t^E + \mathcal{A}_t' \frac{d \log \mathbf{P}_t}{d \log \tau_t} = \mathcal{L}_t' \alpha_t^E. \quad (96)$$

Under the planner's marginal cost pricing, a final good's permit price pass-through $d \log P_{nt} / d \log \tau_t$ is just its embodied emissions cost share. Combining equations (95) and (96), total permit payments are pass-through-weighted final spending, $\tau_t E_t = (d \log \mathbf{P}_t / d \log \tau_t)' \mathbf{V} \mathbf{C}_t^F$.

Next, differentiate equation (94):

$$d \mathbf{S}_t = \mathcal{L}_t [d \mathbf{V} \mathbf{C}_t^F + (d \mathcal{A}_t) \mathbf{S}_t], \quad (97)$$

and differentiate equation (95), substituting equation (97) for $d \mathbf{S}_t$ and $(\alpha_t^E)' \mathcal{L}_t = (d \log \mathbf{P}_t / d \log \tau_t)'$ from equation (96):

$$d E_t = -E_t d \log \tau_t + \frac{1}{\tau_t} \left[(d \alpha_t^E)' \mathbf{S}_t + \left(\frac{d \log \mathbf{P}_t}{d \log \tau_t} \right)' d \mathbf{V} \mathbf{C}_t^F + \left(\frac{d \log \mathbf{P}_t}{d \log \tau_t} \right)' (d \mathcal{A}_t) \mathbf{S}_t \right]. \quad (98)$$

With fixed Cobb-Douglas expenditure shares:

$$d \mathcal{A}_t = \Omega_t \text{diag}(d \alpha_t^M).$$

The first term in equation (98) is the mechanical effect of the higher price at unchanged spending.

Inside the brackets are three effects. The first term is the change in direct emissions intensities at given sales, the second is the reallocation of final demand across final goods, and the third is the recomposition of the input-output network.

The economy-wide elasticity of substitution and MAC elasticity The planner's emissions response is $d \log E_t / d \log \tau_t |_{\bar{Q}_t^F} = -\bar{\sigma}_t(1 - \Gamma_t)$. Inverting this response delivers the MAC elasticity. We build the response from the differentials above and then express the elasticity of substitution $\bar{\sigma}_t$ in terms of observables.

Let $VC_t^F = \sum_n VC_{nt}^F = G_t(\bar{Q}_t^F, \tau_t)$ denote total final spending, and let $w_{n|t} = VC_{nt}^F / VC_t^F$ be the sectoral cost and final spending weights. Define the economy-wide embodied emissions share by:

$$\Gamma_t \equiv \frac{\tau_t E_t}{VC_t^F} = \sum_n w_{n|t} \frac{d \log P_{nt}}{d \log \tau_t}.$$

The second equality divides the pass-through-weighted spending identity $\tau_t E_t = (d \log \mathbf{P}_t / d \log \tau_t)' \mathbf{V} \mathbf{C}_t^F$, from equations (95) and (96), by total final spending. At fixed aggregate final demand, household CES demand gives:

$$\frac{dQ_{nt}^F}{d \log \tau_t} = -v Q_{nt}^F \left(\frac{d \log P_{nt}}{d \log \tau_t} - \Gamma_t \right), \quad \frac{dVC_{nt}^F}{d \log \tau_t} = VC_{nt}^F \left[(1 - v) \frac{d \log P_{nt}}{d \log \tau_t} + v \Gamma_t \right]. \quad (99)$$

Substituting equation (99) into the middle term of equation (98) gives $(d \log \mathbf{P}_t / d \log \tau_t)' d \mathbf{V} \mathbf{C}_t^F = [\Gamma_t^2 + (1 - v) \text{Var}_w(d \log P_{nt} / d \log \tau_t)] VC_t^F d \log \tau_t$. Dividing by $\tau_t E_t = \Gamma_t VC_t^F$ and collecting terms gives:

$$-\frac{d \log E_t}{d \log \tau_t} = (1 - \Gamma_t) + (v - 1) \frac{\text{Var}_w \left(\frac{d \log P_{nt}}{d \log \tau_t} \right)}{\Gamma_t} - \frac{\left(\frac{d \alpha_t^E}{d \log \tau_t} \right)' \mathbf{S}_t + \left(\frac{d \log \mathbf{P}_t}{d \log \tau_t} \right)' \frac{d \mathbf{A}_t}{d \log \tau_t} \mathbf{S}_t}{\tau_t E_t}. \quad (100)$$

Next we define the economy-wide χ_t and $\bar{\sigma}_t$. The economy-wide composite in the elasticity is in embodied terms so materials net out. The composite is then just the economy's capital and labor, with value equal to the non-emissions share of final spending: $\tilde{P}_t^X \tilde{X}_t = (1 - \Gamma_t) VC_t^F$. Since primary factor prices are fixed, the embodied composite price index $\tilde{P}_t^X$ is fixed as well. This gives us that the relative price $\tau_t / \tilde{P}_t^X$ moves one for one with τ_t and changes in composite spending are entirely quantity changes, $d \log \tilde{X}_t = d \log [(1 - \Gamma_t) VC_t^F]$. χ_t and $\bar{\sigma}_t$ are defined as:

$$\chi_t \equiv \frac{\text{Var}_w \left(\frac{d \log P_{nt}}{d \log \tau_t} \right)}{\Gamma_t (1 - \Gamma_t)}, \quad \bar{\sigma}_t \equiv - \frac{d \log (E_t / \tilde{X}_t)}{d \log (\tau_t / \tilde{P}_t^X)} \Big|_{\bar{Q}_t^F}.$$

The expenditure identities $\tau_t E_t = \Gamma_t VC_t^F$ and $\tilde{P}_t^X \tilde{X}_t = (1 - \Gamma_t) VC_t^F$ give $d \log (E_t / \tilde{X}_t) = d \Gamma_t / [\Gamma_t (1 - \Gamma_t)] - d \log \tau_t$. Equation (93) with $VC_t^F = G_t$ gives $d \log VC_t^F / d \log \tau_t = \tau_t E_t / VC_t^F = \Gamma_t$, so

$d \log E_t / d \log \tau_t = (d \Gamma_t / d \log \tau_t) / \Gamma_t - (1 - \Gamma_t)$. Combining the two:

$$\left. \frac{d \log E_t}{d \log \tau_t} \right|_{\bar{Q}_t^F} = -\bar{\sigma}_t(1 - \Gamma_t). \quad (101)$$

Substituting equation (101) into equation (100) and dividing by $1 - \Gamma_t$ gives the exact economy-wide elasticity:

$$\bar{\sigma}_t = 1 + (v - 1)\chi_t - \frac{\left(\frac{d\alpha_t^E}{d \log \tau_t}\right)' \mathbf{S}_t}{\tau_t E_t(1 - \Gamma_t)} - \frac{\left(\frac{d \log \mathbf{P}_t}{d \log \tau_t}\right)' \boldsymbol{\Omega}_t \text{diag}\left(\frac{d\alpha_t^M}{d \log \tau_t}\right) \mathbf{S}_t}{\tau_t E_t(1 - \Gamma_t)}. \quad (102)$$

The last term in equation (102) captures input-output network recomposition. Notice we can rewrite it as:

$$-\frac{\left(\frac{d \log \mathbf{P}_t}{d \log \tau_t}\right)' \boldsymbol{\Omega}_t \text{diag}\left(\frac{d\alpha_t^M}{d \log \tau_t}\right) \mathbf{S}_t}{\tau_t E_t(1 - \Gamma_t)} = -\frac{\sum_n S_{nt} \frac{d\alpha_{nt}^M}{d \log \tau_t} \frac{d \log P_{nt}^M}{d \log \tau_t}}{\tau_t E_t(1 - \Gamma_t)},$$

where $d \log P_{nt}^M / d \log \tau_t = \sum_m \Omega_{mnt} d \log P_{mt} / d \log \tau_t$ is the pass-through of the emissions price into n 's materials bundle, the log derivative of its Cobb-Douglas materials price index.⁵⁶ As the permit price rises, each sector adjusts its materials spending per dollar of sales. By equation (96) the pass-through into its materials bundle is also the embodied emissions cost share of that spending. When sectors cut materials spending where embodied content is high, the input-output network recomposes toward cleaner supply chains.

Next we re-express $\frac{d\alpha_t^E}{d \log \tau_t}$ in equation (102) recalling that we can write $\alpha_{nt}^E = \tau_t(E_{nt}/Y_{nt})/P_{nt}$. For sector n , let $\Gamma_{nt}^G \equiv \sum_i w_{i|n,t} s_{E,int}$ and $1 - \iota_{nt}^{X,G} \equiv \sum_i \omega_{i|n,t}^X (1 - \iota_{int}^X)$ denote the cost-weighted mean of firm emissions cost shares and the composite-cost-weighted mean of firm relative price pass-throughs from equation (62). The intensity response of sector n is just the sector result in equation (66): $-\frac{d \log(E_{nt}/Y_{nt})}{d \log \tau_t} \big|_{Y_{nt}} = \bar{\sigma}_{nt}^{\text{exact}}(1 - \Gamma_{nt}^G)(1 - \iota_{nt}^{X,G})$. Under the planner's marginal cost pricing, the sector pass-through satisfies $1 - d \log P_{nt} / d \log \tau_t = (1 - \Gamma_{nt}^G)(1 - \iota_{nt}^{X,G})$.⁵⁷ The log response is then 1 from the permit price, minus the intensity response, minus the price response: $d \log \alpha_{nt}^E / d \log \tau_t = 1 - \bar{\sigma}_{nt}^{\text{exact}}(1 - d \log P_{nt} / d \log \tau_t) - d \log P_{nt} / d \log \tau_t$. Collecting the $1 - d \log P_{nt} / d \log \tau_t$ terms and multiplying by α_{nt}^E gives:

$$\frac{d\alpha_{nt}^E}{d \log \tau_t} = \alpha_{nt}^E \left(1 - \frac{d \log P_{nt}}{d \log \tau_t}\right) (1 - \bar{\sigma}_{nt}^{\text{exact}}). \quad (103)$$

Substituting equation (103) into equation (102) and using $\alpha_{nt}^E S_{nt} = \tau_t E_{nt}$ gives us:

$$\bar{\sigma}_t = 1 + (v - 1)\chi_t + \frac{\sum_n \frac{E_{nt}}{E_t} (\bar{\sigma}_{nt}^{\text{exact}} - 1) \left(1 - \frac{d \log P_{nt}}{d \log \tau_t}\right)}{1 - \Gamma_t} - \frac{\sum_n S_{nt} \frac{d\alpha_{nt}^M}{d \log \tau_t} \frac{d \log P_{nt}^M}{d \log \tau_t}}{\tau_t E_t(1 - \Gamma_t)}. \quad (104)$$

The materials share derivative can also be re-expressed. Let $s_{M|X,int}$ be firm i 's materials share

⁵⁶This differs from ι_{int}^X which captured pass-through into the productive composite.

⁵⁷Each firm's unit cost satisfies $d \log c_{int} / d \log \tau_t = s_{E,int} + (1 - s_{E,int})\iota_{int}^X$ by Shephard's lemma. Under marginal cost pricing the sector price index aggregates these with revenue weights equal to cost weights $w_{i|n,t}$, and the composite cost weights behind $\iota_{nt}^{X,G}$ make the aggregation factor exactly $d \log P_{nt} / d \log \tau_t = \Gamma_{nt}^G + (1 - \Gamma_{nt}^G)\iota_{nt}^{X,G}$.

within the inner X nest and $s_{M,int} = (1 - s_{E,int})s_{M|X,int}$ its materials share of variable cost. Under the planner's marginal cost pricing, the sector materials share of sales $\alpha_{nt}^M = \sum_i w_{i|n,t} s_{M,int}$ is the cost-weighted mean of the firm shares. We build its derivative in three steps.

First, the firm materials share. Since $s_{M,int} = (1 - s_{E,int})s_{M|X,int}$, its log response splits into: $d \log s_{M,int} = d \log(1 - s_{E,int}) + d \log s_{M|X,int}$. Within the inner production nest, $s_{M|X,int} \propto (P_{nt}^M / P_{int}^X)^{1-\xi_n}$, so $d \log s_{M|X,int} / d \log \tau_t = (1 - \xi_n)(d \log P_{nt}^M / d \log \tau_t - \iota_{int}^X)$ and materials become more expensive relative to the composite by $d \log P_{nt}^M / d \log \tau_t - \iota_{int}^X$, and the share responds with sign $1 - \xi_n$. In the outer nest, equation (55) with the chain rule $d \log(\tau_t / P_{int}^X) / d \log \tau_t = 1 - \iota_{int}^X$ gives $ds_{E,int} / d \log \tau_t = (1 - \sigma_n)s_{E,int}(1 - s_{E,int})(1 - \iota_{int}^X)$, and dividing by $-(1 - s_{E,int})$ turns this into $d \log(1 - s_{E,int}) / d \log \tau_t = -(1 - \sigma_n)s_{E,int}(1 - \iota_{int}^X)$. Adding the two responses and multiplying by $s_{M,int}$:

$$\frac{ds_{M,int}}{d \log \tau_t} = s_{M,int} \left[(1 - \xi_n) \left(\frac{d \log P_{nt}^M}{d \log \tau_t} - \iota_{int}^X \right) - (1 - \sigma_n)s_{E,int}(1 - \iota_{int}^X) \right].$$

Second, the cost share weight. Under marginal cost pricing cost shares are revenue shares, and within-sector CES demand gives $w_{i|n,t} = \phi_{in}(P_{int}/P_{nt})^{1-\epsilon_n}$. Taking logs and differentiating:

$$\frac{dw_{i|n,t}}{d \log \tau_t} = (1 - \epsilon_n)w_{i|n,t} \left(\frac{d \log P_{int}}{d \log \tau_t} - \frac{d \log P_{nt}}{d \log \tau_t} \right). \quad (105)$$

Third, combine the two. Differentiating $\alpha_{nt}^M = \sum_i w_{i|n,t} s_{M,int}$ yields:

$$\frac{d\alpha_{nt}^M}{d \log \tau_t} = \sum_i w_{i|n,t} \frac{ds_{M,int}}{d \log \tau_t} + \sum_i s_{M,int} \frac{dw_{i|n,t}}{d \log \tau_t}.$$

Substituting equation (105) into the second sum and using that the sector price index aggregates firm prices with the same weights, $d \log P_{nt} / d \log \tau_t = \sum_i w_{i|n,t} d \log P_{int} / d \log \tau_t$ gives us that the sum is a covariance: $\sum_i s_{M,int}(1 - \epsilon_n)w_{i|n,t}(d \log P_{int} / d \log \tau_t - d \log P_{nt} / d \log \tau_t) = (1 - \epsilon_n) \text{Cov}_w(s_{M,int}, d \log P_{int} / d \log \tau_t)$. Therefore:

$$\frac{d\alpha_{nt}^M}{d \log \tau_t} = \sum_i w_{i|n,t} \frac{ds_{M,int}}{d \log \tau_t} + (1 - \epsilon_n) \text{Cov}_w \left(s_{M,int}, \frac{d \log P_{int}}{d \log \tau_t} \right).$$

This closes the elasticity: every derivative in equation (104) is now expressed in measured shares and the pass-throughs from equation (96).

The remaining step is to represent the elasticity in the same form as the main text's equation (13) with within and across components. We re-express equation (102) in that form in three steps: (1) define each final good's within response, (2) compute it from the network, and (3) average across final goods. First, by Shephard's lemma, the emissions embodied in one unit of final delivery n are:

$$\tilde{\epsilon}_{nt} = \frac{\partial P_{nt}}{\partial \tau_t} = \frac{P_{nt}}{\tau_t} \frac{d \log P_{nt}}{d \log \tau_t}. \quad (106)$$

The log response of $\tilde{e}_{nt}$ is $d \log \tilde{e}_{nt} / d \log \tau_t = d \log P_{nt} / d \log \tau_t - 1 + d \log(d \log P_{nt} / d \log \tau_t) / d \log \tau_t$, so the within-good decline requires the response of the pass-through itself. Second, differentiating equation (96) gives that response:

$$\frac{d}{d \log \tau_t} \left(\frac{d \log \mathbf{P}_t}{d \log \tau_t} \right) = \mathcal{L}'_t \left[\frac{d \boldsymbol{\alpha}_t^E}{d \log \tau_t} + \frac{d \mathcal{A}'_t}{d \log \tau_t} \frac{d \log \mathbf{P}_t}{d \log \tau_t} \right]. \quad (107)$$

Third, define the within-good elasticity as the average of the final-good responses $-d \log \tilde{e}_{nt} / d \log \tau_t$, weighted by $\frac{d \log P_{nt}}{d \log \tau_t} \left(1 - \frac{d \log P_{nt}}{d \log \tau_t} \right)$:

$$\bar{\sigma}_t^N \equiv - \frac{\sum_n w_{n|t} \frac{d \log P_{nt}}{d \log \tau_t} \frac{d \log \tilde{e}_{nt}}{d \log \tau_t}}{\sum_n w_{n|t} \frac{d \log P_{nt}}{d \log \tau_t} \left(1 - \frac{d \log P_{nt}}{d \log \tau_t} \right)}.$$

Substituting equations (106) and (107) into the numerator and using $\mathcal{L}_t \mathbf{w}_t = \mathbf{S}_t / V C_t^F$ from equation (94) gives:

$$\bar{\sigma}_t^N = 1 - \frac{\left(\frac{d \boldsymbol{\alpha}_t^E}{d \log \tau_t} \right)' \mathbf{S}_t + \left(\frac{d \log \mathbf{P}_t}{d \log \tau_t} \right)' \frac{d \mathcal{A}_t}{d \log \tau_t} \mathbf{S}_t}{\tau_t E_t (1 - \Gamma_t) (1 - \chi_t)}. \quad (108)$$

The weights in the denominator sum to $\Gamma_t (1 - \Gamma_t) - \text{Var}_w \left(\frac{d \log P_{nt}}{d \log \tau_t} \right) = \Gamma_t (1 - \Gamma_t) (1 - \chi_t)$. Multiplying equation (108) by $1 - \chi_t$ and adding $\chi_t v$ then reproduces the right side of equation (102), so:

$$\bar{\sigma}_t = (1 - \chi_t) \bar{\sigma}_t^N + \chi_t v. \quad (109)$$


Equation (109) is the convex combination reported in the main text as equation (13). The weights differ from their sector-level counterparts, where the corresponding coefficient is $\zeta_{nt} = 1 - \chi_{nt} + \Delta_{nt}^X$. The pass-through correction Δ_{nt}^X arises at the sector level because a firm's emissions cost share $s_{E,int}$ and its relative price pass-through $1 - \iota_{int}^X$ are distinct characteristics that vary independently across firms. Economy-wide, the analogous second characteristic is the residual pass-through of a final good's delivery cost beyond its embodied emissions share. Under marginal cost pricing the non-emissions component of embodied cost is capital and labor at fixed prices, so the pass-through $d \log P_{nt} / d \log \tau_t$ equals the embodied share exactly and the residual is zero for every final good. The covariance determining the pass-through correction is zero while the dispersion of the shares themselves survives as χ_t . It does appear in the decentralized schedule above because markups separate embodied shares from pass-through.

Inverting equation (101), the MAC slope is:

$$\left. \frac{d \log MAC_t}{d \log E_t} \right|_{\bar{Q}_t^F} = - \frac{1}{\bar{\sigma}_t (1 - \Gamma_t)}. \quad (110)$$

Under marginal cost pricing, $\mathbf{D}_\mu = \mathbf{I}$ and $\tilde{v}_{nt}^X = 0$, so $\Delta_t^\mu = 0$ and the decentralized system reduces to equations (101) and (110). This is the competitive implementation of the planner.

Figure A19: The Economy-Wide Planner MAC and the Decentralized Schedule



Notes: This figure plots the economy-wide planner MAC curve against the decentralized price-emissions schedule, each constructed as the baseline-revenue-weighted geometric mean of year-specific curves across 2005–2021. Aggregate household final demand is held fixed along each curve, and gross production follows from the input-output network. The red dashed line traces emissions against the permit price in the decentralized equilibrium under observed markups. The blue solid line traces the same economy under marginal cost pricing with unchanged demand elasticities, the competitive implementation whose price-emissions schedule coincides with the planner MAC.

H.3.6 The Decentralized Versus Planner MAC

We compare the entire planner MAC curve with the decentralized schedule by solving the model twice, under observed markups and under marginal cost pricing, with unchanged demand elasticities. Using the same revenue weighting as the main text, we find the competitive allocation emits 37% more at the observed permit price. The primary reason is that markups distort the cost of intermediate inputs. When they are removed, intermediate use and gross production expand. Greater production then carries more emissions. Figure A19 plots the two traced curves. The planner MAC lies above the decentralized schedule at every common emissions level: reaching a given level of emissions costs, on the margin, 79–135% more in resources than the decentralized permit price suggests. Note that the partial equilibrium structure plays a role in the gap. Capital and labor prices are fixed and factor supplies are unconstrained. Removing markups expands intermediate use and shifts factor demand without an offsetting factor price response.

I Microfoundation for the Clean Technology Local Projection

This appendix provides more context around an investment rule and its connection to the local projection specification used in Section 3.4. Moreover, this section discusses how our approach connects to the literature on directed technical change. Throughout the section, we take the approach of utilizing a reduced-form investment rule rather than formulating and solving a full

dynamic investment problem. Our approach amounts to taking a first-order approximation to a general dynamic investment policy function. We supplement this approximation with three additional assumptions: (i) the firm’s forecasts of future prices are linear in current log prices; (ii) investments in clean technology require “time-to-build” that delays their impact by a single (annual) period, so that the investments chosen at date $t - 1$ only become available at date t ; and (iii) there is no depreciation of the clean technology stock, $\delta_B = 0$. Taken together, these assumptions support both the local projection estimating equation and the interpretation of its estimates as the impulse response of clean technology to a change in the carbon price.

I.1 Law of Motion with Time-to-Build Lag and No Depreciation

Let $I_{in,t-1}^B$ denote firm i ’s investment in clean technology undertaken at date $t - 1$. Under the “time-to-build” assumption, this investment becomes productive one period later. So, the law of motion for the level of clean technology, B_{int} , is:

$$B_{int} = B_{in,t-1} + I_{in,t-1}^B.$$

Defining $\iota_{in,t-1}^B \equiv I_{in,t-1}^B / B_{in,t-1}$ as the clean technology investment rate in period $t - 1$ allows us to characterize the law of motion for the log level of clean technology as $\log B_{int} = \log B_{in,t-1} + \log(1 + \iota_{in,t-1}^B)$. Linearizing $\log(1 + \iota)$ around a sector-level steady-state investment rate $\bar{\iota}_n$ gives:

$$\log B_{int} \approx \log B_{in,t-1} + c_n + \kappa_n \iota_{in,t-1}^B, \quad (111)$$

where $\kappa_n = 1/(1 + \bar{\iota}_n)$ and $c_n = \log(1 + \bar{\iota}_n) - \bar{\iota}_n/(1 + \bar{\iota}_n)$ are linearization constants and firm-level heterogeneity in steady-state investment rates is absorbed into the firm-specific intercepts introduced below. The main text law of motion in equation (14) is equation (111) with $c_n + \kappa_n \iota_{in,t-1}^B$ relabeled as the log growth contribution $I_{in,t-1}^B$.

I.2 An Investment Rule as Directed Technical Change

Our characterization of the investment rule proceeds in three steps: (i) taking a general stance on the form of the investment rule, (ii) taking a log first order approximation of this investment rule, and then (iii) combining these steps with an assumption surrounding the forecasts of future prices. Finally, we outline the additional restrictions sufficient to generate a (log) relative price functional form, and how this connects with the literature on directed technical change.

A general investment policy Whatever dynamic problem the firm solves, its solution is a policy function mapping the date- $(t - 1)$ state into the investment rate. Because clean technology takes a period to install, investment undertaken at $t - 1$ is repaid by a stream of cost savings from date t onward. Naturally, the firm’s forecasts of future prices will affect the optimal amount of investment. Writing $p \equiv \log P$ and letting $\mathbb{E}_{t-1}$ denote the expectation conditional on date- $(t - 1)$

information,

$$\iota_{in,t-1}^B = G_n \left(p_{t-1}^E, p_{in,t-1}^X, \{ \mathbb{E}_{t-1} p_{t-1+j}^E \}_{j \geq 1}, \{ \mathbb{E}_{t-1} p_{in,t-1+j}^X \}_{j \geq 1} \right) + \nu_{in,t-1}, \quad (112)$$

where $\nu_{in,t-1}$ is the remaining error term given the arguments, and we have assumed that the arguments of $G_n(\cdot)$ involve the conditional expectation of log prices.

While equation (112) remains quite general, it does omit several factors that the literature on directed technical change has outlined, including the firm's existing capital stock, its own scale, and any aggregate states. To the extent that some of these are time-invariant, they can be interpreted as being absorbed by the firm-level fixed effects introduced in the empirical specification.

A first-order approximation Taking a first order approximation of G_n around a sector-level steady state $(\bar{p}_n^E, \bar{p}_n^X)$ gives:

$$\begin{aligned} \iota_{in,t-1}^B = & \bar{\iota}_n + a_{0,n}(p_{t-1}^E - \bar{p}_n^E) + b_{0,n}(p_{in,t-1}^X - \bar{p}_n^X) + \nu_{in,t-1} \\ & + \sum_{j \geq 1} \left[a_{j,n}(\mathbb{E}_{t-1} p_{t-1+j}^E - \bar{p}_n^E) + b_{j,n}(\mathbb{E}_{t-1} p_{in,t-1+j}^X - \bar{p}_n^X) \right], \end{aligned} \quad (113)$$

where $a_{j,n}$ and $b_{j,n}$ are the partial effects of the date- $(t-1+j)$ price forecasts on current investment.

Linear price forecasts We assume the two date- $(t-1)$ prices span the firm's forecast at every horizon:

$$\begin{aligned} \mathbb{E}_{t-1} p_{t-1+j}^E &= \mu_{j,n}^E + \rho_{j,n}^{EE} p_{t-1}^E + \rho_{j,n}^{EX} p_{in,t-1}^X, \\ \mathbb{E}_{t-1} p_{in,t-1+j}^X &= \mu_{j,n}^X + \rho_{j,n}^{XE} p_{t-1}^E + \rho_{j,n}^{XX} p_{in,t-1}^X. \end{aligned} \quad (114)$$

If this assumed form holds, every horizon- j forecast is a linear function of p_{t-1}^E and $p_{in,t-1}^X$. Any weighted sum of forecasts therefore collapses into a constant plus loadings on those two prices, which is why no separate expectation term survives in the estimated rule.

Substituting equation (114) into equation (113) and collecting terms yields a rule in date- $(t-1)$ prices alone:

$$\iota_{in,t-1}^B = \phi_{0,in} + \phi_n^E p_{t-1}^E + \phi_n^X p_{in,t-1}^X + \nu_{in,t-1}, \quad (115)$$

with:

$$\phi_n^E = a_{0,n} + \sum_{j \geq 1} (a_{j,n} \rho_{j,n}^{EE} + b_{j,n} \rho_{j,n}^{XE}), \quad \phi_n^X = b_{0,n} + \sum_{j \geq 1} (a_{j,n} \rho_{j,n}^{EX} + b_{j,n} \rho_{j,n}^{XX}), \quad (116)$$

and $\phi_{0,in}$ collecting the steady-state investment rate, the forecast intercepts, and firm-specific scale. Equation (115) is the main text rule (16), with $\beta_n^E = \kappa_n \phi_n^E$ and $\beta_n^X = \kappa_n \phi_n^X$ for the linearization constant κ_n of equation (111).

Equation (116) suggests limits on how the resulting estimates should be interpreted. First,

the coefficients of interest in the local projection are best interpreted as reduced form coefficients themselves. Even in the derivation presented here, the estimates from the local projections represent a combination of the effects of both an investment policy rule and the process that generates the evolution of prices (and their forecasts). Therefore, our counterfactuals should be interpreted as requiring that both of these dynamic aspects of the environment be held constant.

The relative-price restriction While not strictly necessary for the local projections, two further conditions reduce equation (115) to a rule in the relative price alone. The first is that G_n is homogeneous of degree zero in prices, so that scaling all current and expected future prices by a common factor leaves the investment rate unchanged:

$$a_{0,n} + b_{0,n} + \sum_{j \geq 1} (a_{j,n} + b_{j,n}) = 0.$$

The second is that forecasts inherit the same scaling, $\rho_{j,n}^{EE} + \rho_{j,n}^{EX} = 1$ and $\rho_{j,n}^{XE} + \rho_{j,n}^{XX} = 1$ at every horizon. Applying both to (116) gives $\phi_n^E + \phi_n^X = 0$. Writing $\phi_n^{EX} \equiv \phi_n^E = -\phi_n^X$, the investment rule becomes:

$$\iota_{in,t-1}^B = \phi_{0,in} + \phi_n^{EX} \log \left(\frac{P_{t-1}^E}{P_{in,t-1}^X} \right) + \nu_{in,t-1}, \quad \phi_n^{EX} > 0, \quad (117)$$

which is the form we carry through the rest of this appendix. It remains a local approximation around the steady state, and presumes $\iota_{in,t-1}^B > -1$ so that the log law of motion is well defined.

The sign restriction and directed technical change The functional form of equation (117) follows from the approximation above, while a prediction of its sign is connected to the theory of directed technical change. In some of that literature, the return to factor-augmenting innovation is proportional to the factor's total expenditure rather than to its price alone. For instance, Acemoglu (2002) decomposes the relative profitability of innovating for one factor rather than another into a price effect and a market size effect, while Acemoglu et al. (2012) provide the clean-versus-dirty analogue, where the ratio of expected profits from clean and dirty research is the product of a price effect, a market size effect, and a direct productivity effect.

Our setting collapses the first two of these into a single regressor. By Shephard's lemma the marginal return to clean technology is the emissions cost share, $\partial \log MC_{int} / \partial \log B_{int} = -s_{E,int}$, so the directed technical change incentive is $s_{E,int}$ scaled by firm revenue. The two-input outer nest gives $\partial \log s_{E,int} / \partial \log (P_t^E / P_{int}^X) = (1 - \sigma_n)(1 - s_{E,int})$, so the emissions cost share rises with the relative price of emissions if and only if $\sigma_n < 1$. The price and market size effects are therefore not separately identified here: they move together, and $\log(P_{t-1}^E / P_{in,t-1}^X)$ is a first-order sufficient statistic for both. This is the gross complements case that Acemoglu et al. (2012) describe, in which the price and market size effects reinforce one another rather than offset. Because we estimate σ_n close to zero in every sector, it is the relevant case for our application, and it delivers $\phi_n^{EX} > 0$ as a prediction of the theory given our estimated σ_n rather than as an assumption of convenience. The

restriction this leaves in place is that firm revenue, which scales the market size effect, sits in $\phi_{0,in}$ and is treated as fixed within firm.

I.3 Combined Structural Law of Motion

Substituting equation (117) into equation (111) and defining $\beta_n^{EX} \equiv \kappa_n \phi_n^{EX}$, the structural law of motion for $\log B_{int}$ is:

$$\log B_{int} = \alpha_{in} + \log B_{in,t-1} + \beta_n^{EX} \log \left(\frac{P_{t-1}^E}{P_{in,t-1}^X} \right) + u_{int}, \quad (118)$$

where $\alpha_{in} \equiv c_n + \kappa_n \phi_{0,in}$ absorbs the linearization constant and firm-specific steady-state drift and $u_{int} \equiv \kappa_n \nu_{in,t-1}$. In first differences,

$$\Delta \log B_{int} = \alpha_{in} + \beta_n^{EX} \log \left(\frac{P_{t-1}^E}{P_{in,t-1}^X} \right) + u_{int}. \quad (119)$$

The main text local projection in equation (17) is obtained by iterating equation (119) forward and projecting on date- $(t-1)$ prices. The forecasting assumption in equation (114) makes the resulting two-regressor form exact, together with two further conditions: (i) the forecasting coefficients are common across firms within a sector, and (ii) the date- $(t-1)$ residuals $\nu_{in,t-1}$ are orthogonal to the two date- $(t-1)$ prices. Under these conditions the coefficients may either impose the $\beta_n^E = -\beta_n^X$ restriction of equation (117) or load independently as in equation (115).

I.4 Moving Average Representation and Identification of σ_n

Iterating (118) backward from an initial date t_0 yields:

$$\log B_{int} = \log B_{in,t_0} + (t - t_0)\alpha_{in} + \beta_n^{EX} \sum_{s=1}^{t-t_0} \log \left(\frac{P_{t-s}^E}{P_{in,t-s}^X} \right) + \tilde{u}_{int}, \quad (120)$$

with $\tilde{u}_{int} \equiv \sum_{\tau=t_0+1}^t u_{in\tau}$. Because $\delta_B = 0$, past relative prices enter with equal unit weights.

Substituting equation (120) into the static first-order condition in equation (5) delivers the reduced-form specification for the emissions-to-composite ratio when B_{int} is endogenous. Crucially, B_{int} in equation (120) depends only on information dated $t-1$ or earlier, so the surprise satisfies $\mathbb{E}[z_t \log B_{int}] = 0$. The endogenous directed technical change channel cannot contaminate the within-period response because B_{int} is predetermined with respect to period- t prices, and no control for the history of relative prices is needed for this to hold.

I.5 Local Projection Coefficient and Carbon Price Surprise Orthogonality

Let z_{t-1} denote the carbon price surprise realized in period $t-1$; it is the instrument for $\log P_{t-1}^E$ in the empirical local projection (21), whose outcome is differenced from $t-1$ to $t+h$. Two conditions

ensure that the necessary orthogonality condition holds for the local projections. First, the surprise is unforecastable, and is therefore orthogonal to all period- $(t - 2)$ and earlier firm and aggregate information, so

$$\mathbb{E}[z_{t-1} \tilde{u}_{in,t-1}] = 0. \quad (121)$$

Second, the surprise moves investment only through relative prices, which excludes the possibility that the investment residuals respond to the surprise at any horizon,

$$\mathbb{E}[z_{t-1} \nu_{in,s}] = 0 \quad \text{for all } s \geq t - 1. \quad (122)$$

Condition (121) is supported by the high-frequency construction of the surprise series. Condition (122) is the exclusion restriction that is more sensitive to our modeling approach and the assumptions outlined in this section, and requires that regulatory news shift clean technology investment only by moving the relative factor prices that enter the investment rule.

By equation (120), the outcome window contains the relative prices dated $t - 1$ through $t + h - 1$ and the innovations $u_{in,t}, \dots, u_{in,t+h}$:

$$\log B_{in,t+h} - \log B_{in,t-1} = (h + 1) \alpha_{in} + \beta_n^{EX} \sum_{\ell=0}^h \log \left(\frac{P_{t-1+\ell}^E}{P_{in,t-1+\ell}^X} \right) + \sum_{\tau=t}^{t+h} u_{in\tau}. \quad (123)$$

When equation (123) is projected on z_{t-1} , condition (122) removes every innovation in the window (each $u_{in\tau} = \kappa_n \nu_{in,\tau-1}$ is dated $t - 1$ or later), and the price terms load through their first-stage responses. Let

$$\psi_{\ell,n} \equiv \frac{\partial \mathbb{E} \left[\log \left(P_{t-1+\ell}^E / P_{in,t-1+\ell}^X \right) \right]}{\partial z_{t-1}}, \quad \ell = 0, 1, \dots,$$

denote the first-stage response of the log relative price at horizon ℓ to the surprise, with controls partialled out as in the estimating equation. These responses involve the firm-level price $P_{in,t-1+\ell}^X$; writing them with a sector index presumes a common within-sector response, and with heterogeneous firm responses, $\psi_{\ell,n}$ is the corresponding IV-weighted sector average. The reduced-form projection coefficient is $R_{h,n} \equiv \beta_n^{EX} \sum_{\ell=0}^h \psi_{\ell,n}$, in units of the surprise. The empirical specification (21) instead reports an IV coefficient on $\log P_{t-1}^E$, which rescales the reduced form by the impact first stage $\psi_0^E \equiv \partial \mathbb{E}[\log P_{t-1}^E] / \partial z_{t-1}$. The model-implied local projection coefficient is therefore:

$$\Psi_{h,n}^E = \frac{\beta_n^{EX}}{\psi_0^E} \sum_{\ell=0}^h \psi_{\ell,n}. \quad (124)$$

The sum has $h + 1$ terms, matching the $h + 1$ investment periods $I_{in,t-1}^B, \dots, I_{in,t+h-1}^B$ inside the outcome window. In particular the impact coefficient is nonzero: at $h = 0$ the outcome is $\log B_{int} - \log B_{in,t-1} = \alpha_{in} + \beta_n^{EX} \log(P_{t-1}^E / P_{in,t-1}^X) + u_{int}$, so $\Psi_{0,n}^E = \beta_n^{EX} \psi_{0,n} / \psi_0^E$, the direct effect of the date- $(t - 1)$ relative price on date- $(t - 1)$ investment. Section 4.3 describes the IV local projections that estimate these coefficients through the observable-ratio counterpart derived in the next subsection;

in the just-identified case the IV coefficient equals this reduced-form-to-first-stage ratio.

I.6 Mapping Clean-Technology Dynamics to Observables

Because B_{int} is not directly observed, the local projections in Section 4.3 are estimated on the observable composite-to-emissions ratio rather than on clean technology itself. This subsection derives the mapping between the two.

The static first-order condition in equation (5) implies:

$$\log \left(\frac{X_{int}}{E_{int}} \right) = \kappa_{in} + \sigma_n (\log P_t^E - \log P_{int}^X) + (1 - \sigma_n) \log B_{int}, \quad (125)$$

where $\kappa_{in} \equiv -\sigma_n \log (\pi_{in}/(1 - \pi_{in}))$, so the observable ratio combines a contemporaneous static-substitution response to the emissions-composite price ratio with a $(1 - \sigma_n)$ -weighted loading on clean technology. Under the linear price-forecasting conditions that deliver the two-regressor form of equation (17), taking conditional long differences between $t - 1$ and $t + h$ and substituting that projection yields the observable-ratio local projection:

$$\mathbb{E}_{t-1} \left[\log \frac{X_{in,t+h}}{E_{in,t+h}} - \log \frac{X_{in,t-1}}{E_{in,t-1}} \right] = a_{in}^{(h)} + \gamma_{h,n}^E \log P_{t-1}^E + \gamma_{h,n}^X \log P_{in,t-1}^X, \quad (126)$$

whose coefficients combine the dynamic clean-technology response from equation (17) with the static CES substitution channel from equation (125).

The window algebra of equations (123)–(124) makes the combination explicit for the coefficient on the emissions price in the just-identified IV specification (21). Projecting the long difference of equation (125) on the surprise z_{t-1} —the window innovations dropping by condition (122) as in the previous subsection—gives the reduced-form response $\sigma_n(\psi_{h+1,n} - \psi_{0,n}) + (1 - \sigma_n)\beta_n^{EX} \sum_{\ell=0}^h \psi_{\ell,n}$, where the terminal static term carries index $h + 1$ because the outcome window ends at $t + h = t - 1 + (h + 1)$. Rescaling by the impact first stage ψ_0^E and using equation (124) yields:

$$\gamma_{h,n}^E = (1 - \sigma_n) \Psi_{h,n}^E + \sigma_n \frac{\psi_{h+1,n} - \psi_{0,n}}{\psi_0^E},$$

which decomposes the IV coefficient into a $(1 - \sigma_n)$ -weighted clean technology impulse response and the static substitution induced by the movement in the relative factor price over the outcome window. As $\sigma_n \rightarrow 0$, the static channel shuts down and $\gamma_{h,n}^E \rightarrow \Psi_{h,n}^E$. Since our sector-level estimates of σ_n are all close to zero, the two objects are close in practice.

In principle, permit price shocks can also move $P_{in,t-1}^X$ through input-output linkages, so $\gamma_{h,n}^X$ is best interpreted as a reduced-form control coefficient rather than a stand-alone structural elasticity. Finally, in our implementation of the local projection we measure the composite input quantity in equation (126) as non-emissions spending deflated by the composite price index, $X_{int} = (C_{int}^K + C_{int}^L + C_{int}^M)/P_{int}^X$, as in Section 4.1.2.

J External Benchmark: Translating the Colmer et al. (2024) Treatment Effect

Section 5.3 benchmarks our clean technology elasticity against Colmer et al. (2024). Colmer et al. (2024) estimate that EU-ETS Phase II treatment reduced French firms’ logged emissions intensity, measured as $\log(\text{CO}_2/\text{value added})$, by $\beta^{\text{CMMW}} = -0.174$ (standard error 0.075; their Table 2, column 5). In their empirical design, control firms are not subject to EU-ETS and face a carbon price implicit in pre-existing energy taxation, P_0 . Treated firms originally faced just P_0 , but then moved to $P_0 + \bar{P}^E$ under the ETS. The average ETS price over their 2008–2012 treatment window is $\bar{P}^E = 18.08$ USD/tCO₂. Dividing the treatment effect by the implied log price change gives a price elasticity of emissions intensity: $\Psi^E = -\beta^{\text{CMMW}} / \log((P_0 + \bar{P}^E)/P_0)$. This is essentially the same object as our clean technology local projection estimate, but constructed from their estimates. To compare the different approaches, we need an estimate from our data with the same sector coverage as their sample, and a value for the baseline price P_0 .

J.1 Sector Coverage

Their estimation sample comes from the EACEI energy survey. The survey covers all of French manufacturing except energy production, agri-food, and sawmills. Refining is included under energy activities in the survey, and mining is outside manufacturing altogether. To get the best comparison, we restrict our estimating sample to sectors included in the Colmer et al. (2024) sample. These sectors are basic metals, cement/lime/glass/ceramics, chemicals, fabricated metal products, pulp/paper/board, and rubber and plastic products. Our light manufacturing sector pools the excluded agri-food and sawmill activities with other light manufacturing activities that are included, so we report estimates with and without it. The regressions are identical to our pooled estimates in Appendix A, except for the change of sample. Our estimates are $\hat{\gamma}_4^E = 0.119$ (SE 0.067; $N = 8,171$) without light manufacturing and 0.153 (SE 0.064; $N = 10,912$) with light manufacturing.⁵⁸

J.2 Determining Baseline Prices P_0

The baseline price P_0 is the effective carbon price on industrial energy outside EU-ETS. We place it roughly between 2 and 6.5 EUR/tCO₂ over 2008–2012. Statutory excises on natural gas, coal, and heavy fuel oil corresponded to about 3.5–6.5 EUR/tCO₂ (Marten, 2023). OECD data put the emissions-weighted non-ETS tax on French industrial energy at 4.55 EUR/tCO₂ in 2012 (OECD, 2013, 2016). Colmer et al. (2024) themselves, citing Callonnec (2009), report a nominal rate of about 6 EUR/tCO₂ on fuel oil and natural gas, with industrial coal and a broad range of industrial gas uses fully exempt.

⁵⁸The sectors excluded from their sample, such as power, refining, mining, and other non-manufacturing, give 0.343 (SE 0.094).

Table A17: Benchmarking Against the Colmer et al. (2024) Treatment Effect

<i>Panel A: implied elasticity at documented baseline prices</i>			
Baseline anchor	P_0 (EUR)	P_0 (USD)	Implied Ψ^E [95% CI]
Central range, low	2.00	2.75	0.086 [0.013, 0.158]
Central range	3.00	4.12	0.103 [0.016, 0.191]
OECD sector average	4.55	6.25	0.128 [0.020, 0.236]
Statutory gas excise	6.50	8.93	0.157 [0.024, 0.290]
<i>Panel B: LP-IV samples and reconciling baseline price</i>			
LP-IV sample	$\hat{\gamma}_4^E$ (SE)	P_0^* (EUR)	P_0^* (USD)
Colmer-overlap manufacturing	0.119 (0.067)	3.99	5.49
+ Light Manufacturing	0.153 (0.064)	6.21	8.53
All sectors	0.211 (0.043)	10.29	14.14

Notes: This table benchmarks our clean technology elasticity against Colmer et al. (2024). Panel A translates their Phase II treatment effect on logged emissions intensity (-0.174 , SE 0.075) into a price elasticity $\Psi^E = -\beta^{\text{CMMW}} / \log((P_0 + \bar{P}^E)/P_0)$ at documented values of the baseline non-ETS carbon price P_0 , using the mean Phase II permit price $\bar{P}^E$ of \$18.08. The central range reflects statutory excises net of exemptions (Callonnec, 2009; Marten, 2023). The OECD sector average is the emissions-weighted non-ETS rate on French industrial energy in 2012 (OECD, 2016). The statutory gas excise is the excise on natural gas (OECD, 2013). Euro values convert at the 2008–2012 mean exchange rate of 1.37 USD/EUR. Panel B reports our pooled five-year local projection estimate for each sector scope and the reconciling baseline price P_0^* at which the implied elasticity in Panel A equals our point estimate. The Colmer-overlap scope covers basic metals, cement/lime/glass/ceramics, chemicals, fabricated metal products, pulp/paper/board, and rubber and plastic products. Standard errors in Panel B are Driscoll-Kraay.

J.2.1 Estimated Colmer Elasticities

Table A17 Panel A reports the elasticities implied by the Colmer et al. (2024) estimate across the above range of prices. We convert euros at the 2008–2012 mean exchange rate of 1.37 USD/EUR to be in the same monetary units as our paper. The implied Ψ^E rises from 0.086 at 2 EUR to 0.157 at the statutory gas rate. In Panel B, we instead invert the mapping to determine the baseline price at which the Colmer-implied elasticity equals our point estimate exactly. That price is 3.99 EUR (\$5.49) without light manufacturing and 6.21 EUR (\$8.53) with it. Both are inside the documented range of prices.⁵⁹ Their confidence interval contains our point estimate at any baseline above 0.96 EUR without light manufacturing, and above 1.84 EUR with it. The agreement does not depend on where in the documented range the true baseline falls.

K Long Run Elasticity of Substitution with Endogenous Clean Technology

This appendix derives the long-run elasticity of substitution between emissions and the productive composite when clean technology responds to the permit price. The short-run elasticity σ_n holds clean technology B_{int} fixed. The long-run elasticity allows B_{int} to change with the permit price.

⁵⁹The all-sector pooled estimate requires 10.29 EUR (\$14.14), above the entire range.

The long-run elasticity therefore combines within-period input substitution with induced clean technology adoption.

Consider a permanent shock to the permit price at date $t-1$: banking ties current and expected future permit prices together, so the shock moves the whole price path and $d \log P_{t+h}^E / d \log P_{t-1}^E = 1$ at every horizon. We take the five-year horizon of our local projection estimates as the long run, evaluate all responses there, and suppress time subscripts on long-run outcomes. Every differential below is taken with respect to the date- $(t-1)$ shock. Define the long-run permit price elasticity of clean technology as:

$$\psi_n^B \equiv \frac{d \log B_{in}}{d \log P_{t-1}^E},$$

the total response of clean technology five years after the shock, which we treat as common to firms within a sector. This is the object the five-year local projection coefficient $\Psi_{h,n}^E$ estimates (Appendix I). When $\psi_n^B > 0$, higher permit prices induce firms to adopt cleaner technology. Because the elasticity of substitution is defined with respect to the relative emissions price, we also need the pass-through of the permit price into the firm's productive composite price. We use the within-period pass-through ι_{int}^X of Appendix H.1, treating it as constant across horizons:

$$\iota_{int}^X \equiv \frac{d \log P_{in}^X}{d \log P_{t-1}^E},$$

so the implied response of the log relative emissions price is:

$$\frac{d \log (P^E / P_{in}^X)}{d \log P_{t-1}^E} = 1 - \iota_{int}^X.$$

The long-run micro elasticity is the total response of the emissions-to-composite ratio to the relative emissions price, allowing clean technology to adjust:

$$\sigma_{in}^{LR} \equiv -\frac{d \log (E_{in} / X_{in})}{d \log (P^E / P_{in}^X)}. \quad (127)$$

The short-run first-order condition in equation (5), evaluated in the long run, implies:

$$\log \left(\frac{E_{in}}{X_{in}} \right) = \sigma_n \log \left(\frac{\pi_{in}}{1 - \pi_{in}} \right) - \sigma_n \log \left(\frac{P^E}{P_{in}^X} \right) + (\sigma_n - 1) \log B_{in}. \quad (128)$$

Differentiating equation (128) with respect to the date- $(t-1)$ shock, and allowing B_{in} and P_{in}^X to respond, gives:

$$-\frac{d \log (E_{in} / X_{in})}{d \log P_{t-1}^E} = \sigma_n (1 - \iota_{int}^X) - (\sigma_n - 1) \psi_n^B. \quad (129)$$

Combining equations (127)–(129) gives:

$$\begin{aligned}
\sigma_{in}^{LR} &= -\frac{d \log(E_{in}/X_{in})}{d \log(P^E/P_{in}^X)} \\
&= -\frac{\frac{d \log(E_{in}/X_{in})}{d \log P_{t-1}^E}}{\frac{d \log(P^E/P_{in}^X)}{d \log P_{t-1}^E}} \\
&= -\frac{-\sigma_n(1 - \iota_{int}^X) + (\sigma_n - 1)\psi_n^B}{1 - \iota_{int}^X} \\
&= \sigma_n - (\sigma_n - 1)\frac{\psi_n^B}{1 - \iota_{int}^X} \\
&= \sigma_n + (1 - \sigma_n)\frac{\psi_n^B}{1 - \iota_{int}^X}, \tag{130}
\end{aligned}$$

provided $\iota_{int}^X < 1$. The long-run elasticity equals the short-run within-firm elasticity plus the clean technology response, weighted by $1 - \sigma_n$ and divided by the relative price pass-through $1 - \iota_{int}^X$, the movement in the relative emissions price. Equation (130) makes clear when endogenous clean technology raises the elasticity. If emissions and the productive composite are complements, $\sigma_n < 1$, then a positive clean technology response raises σ_{in}^{LR} above the short-run elasticity. In the empirically relevant case where $\sigma_n \approx 0$, the long-run elasticity is approximately $\psi_n^B/(1 - \iota_{int}^X)$. In our simulations ι_{int}^X is near zero for the median and average firm-year, so the local projection estimates approximately identify the long-run elasticity.

Appendix References

- Acemoglu, Daron (2002) “Directed Technical Change,” *The Review of Economic Studies*, Vol. 69, No. 4, pp. 781–809.
- Andrews, Isaiah, James H Stock, and Liyang Sun (2019) “Weak instruments in instrumental variables regression: Theory and practice,” *Annual Review of Economics*, Vol. 11, No. 1, pp. 727–753.
- Baqaei, David Rezza and Emmanuel Farhi (2020) “Productivity and Misallocation in General Equilibrium,” *The Quarterly Journal of Economics*, Vol. 135, No. 1, pp. 105–163.
- Baumol, William J. and Wallace E. Oates (1988) *The Theory of Environmental Policy*, Cambridge, UK: Cambridge University Press, 2nd edition.
- 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.

- Boyd, Stephen and Lieven Vandenberghe (2004) *Convex Optimization*, Cambridge: Cambridge University Press.
- Callonnec, Gaël (2009) “Fiscalité comparée de l’énergie et du CO₂ en Europe et en France,” *ADEME & vous, Stratégie & études*, No. 20, pp. 1–6.
- Copeland, Brian R. and M. Scott Taylor (1994) “North-South Trade and the Environment,” *Quarterly Journal of Economics*, Vol. 109, No. 3, pp. 755–787.
- (2003) *Trade and the Environment: Theory and Evidence*, Princeton, NJ: Princeton University Press.
- Cropper, Maureen L. and Wallace E. Oates (1992) “Environmental Economics: A Survey,” *Journal of Economic Literature*, Vol. 30, No. 2, pp. 675–740.
- Davis, Steven J. and Ken Caldeira (2010) “Consumption-Based Accounting of CO₂ Emissions,” *Proceedings of the National Academy of Sciences*, Vol. 107, No. 12, pp. 5687–5692.
- Ebert, Udo and Heinz Welsch (2007) “Environmental Emissions and Production Economics: Implications of the Materials Balance,” *American Journal of Agricultural Economics*, Vol. 89, No. 2, pp. 287–293.
- Hulten, Charles R. (1978) “Growth Accounting with Intermediate Inputs,” *The Review of Economic Studies*, Vol. 45, No. 3, pp. 511–518.
- Johnson, Robert C. and Guillermo Noguera (2012) “Accounting for Intermediates: Production Sharing and Trade in Value Added,” *Journal of International Economics*, Vol. 86, No. 2, pp. 224–236.
- León-Ledesma, Miguel A., Peter McAdam, and Alpo Willman (2010) “Identifying the Elasticity of Substitution with Biased Technical Change,” *American Economic Review*, Vol. 100, No. 4, pp. 1330–1357.
- Leontief, Wassily (1970) “Environmental Repercussions and the Economic Structure: An Input-Output Approach,” *The Review of Economics and Statistics*, Vol. 52, No. 3, pp. 262–271.
- Marten, Mélanie (2023) “Carbon Pricing Reform and Expectations: Evidence from French Manufacturing, 2005–2019,” THEMA Working Paper 2022-19, THEMA, CY Cergy Paris Université.
- McFadden, Daniel (1978) “Cost, Revenue, and Profit Functions,” in Melvyn Fuss and Daniel McFadden eds. *Production Economics: A Dual Approach to Theory and Applications, Volume 1*, Amsterdam: North-Holland, pp. 3–109.
- Murty, Sushama (2015) “On the Properties of an Emission-Generating Technology and Its Parametric Representation,” *Economic Theory*, Vol. 60, No. 2, pp. 243–282.

- Murty, Sushama and R. Robert Russell (2018) “Modeling Emission-Generating Technologies: Reconciliation of Axiomatic and By-Production Approaches,” *Empirical Economics*, Vol. 54, No. 1, pp. 7–30.
- Murty, Sushama, R. Robert Russell, and Steven B. Levkoff (2012) “On Modeling Pollution-Generating Technologies,” *Journal of Environmental Economics and Management*, Vol. 64, No. 1, pp. 117–135.
- OECD (2013) *Taxing Energy Use: A Graphical Analysis*, Paris: OECD Publishing.
- (2016) *Effective Carbon Rates: Pricing CO₂ through Taxes and Emissions Trading Systems*, Paris: OECD Publishing.
- Pasinetti, Luigi L. (1973) “The Notion of Vertical Integration in Economic Analysis,” *Metroeconomica*, Vol. 25, No. 1, pp. 1–29.
- Pethig, Rüdiger (1976) “Pollution, Welfare, and Environmental Policy in the Theory of Comparative Advantage,” *Journal of Environmental Economics and Management*, Vol. 2, No. 3, pp. 160–169.
- (2006) “Non-linear Production, Abatement, Pollution and Materials Balance Reconsidered,” *Journal of Environmental Economics and Management*, Vol. 51, No. 2, pp. 185–204.
- Shephard, Ronald W. (1970) *Theory of Cost and Production Functions*, Princeton, NJ: Princeton University Press.
- Siebert, Horst, Jürgen Eichberger, Ralf Gronych, and Rüdiger Pethig (1980) *Trade and Environment: A Theoretical Enquiry*, Amsterdam: Elsevier.