The Spatial Consequences of Industrial Environmental Regulation*

Cheolhwan Kim†

Yale University

November 3, 2026

Click Here for the Latest Version

Abstract

I study how South Korea's place-based cap-and-trade system affects local air pollution and production. Difference-in-differences estimates show that manufacturing emissions fall in regulated districts relative to the same industries in unregulated districts. Plant and district comparisons reveal abatement through electrification and emission-control equipment, as well as patterns consistent with pollution shifts within firms, to unregulated districts, and to power generation. To account for such shifts and translate these responses into aggregate and spatial outcomes, I build a spatial equilibrium model of manufacturing and upstream power generation. The model attributes 19% of the estimated relative emissions decline to geographic pollution shifts. Still, aggregate pollution damage falls by 38% at a consumption cost of 1%, because damage reductions in high-damage districts outweigh increases elsewhere. Holding national emissions totals fixed for each pollutant, a nationwide alternative with damage-proportional emissions prices achieves larger damage reductions at a lower consumption cost than the actual policy.

*I am grateful for the support of the Economic Growth Center at Yale. Seunghak Lee provided outstanding research assistance. I am deeply grateful to Costas Arkolakis, Kenneth Gillingham, Amit Khandelwal, and Samuel Kortum for their encouragement and guidance. I thank Mayara Felix, Teresa Fort, Niclas Moneke, Emmanuel Murray Leclair, Rohini Pande, and Nicholas Ryan for their invaluable feedback and participants at seminars for their helpful comments. All errors are mine. The statistical analysis of firm- and plant-level data in South Korea reported in this study was conducted at a Statistics Data Center of Statistics Korea under arrangements that maintained legal confidentiality requirements (Project Number 5465). Any views expressed are those of the author and not those of the Korean Ministry of Environment or Statistics Korea. Statistics Korea's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this information product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. The author declares that no relevant or material financial interests are related to the research described in this paper.

†Email: cheolhwan.kim@yale.edu

1 Introduction

Local air pollution imposes large health costs that, unlike the climate effects of greenhouse gases, depend on where emissions occur. Permanently meeting the World Health Organization (WHO) guideline for fine particulate matter ($PM_{2.5}$) would increase average life expectancy by an estimated 2.3 years worldwide and 11.9 years in Delhi (Greenstone and Hasenkopf 2023). Regulations therefore often target large polluting manufacturing and power plants in densely populated, heavily polluted regions where pollution damages are high.[1] In response, plants can abate emissions with the use of cleaner technologies or shift pollution elsewhere within and across firms, sectors, and regions.

Such pollution shifts complicate policy evaluation because emissions reductions at regulated units may reflect displacement to unregulated units rather than a decline in aggregate emissions. When pollution shifts raise emissions at unregulated units, comparisons between regulated and unregulated units can overstate aggregate emissions reductions. At the same time, shifting emissions toward lower-damage locations can decrease aggregate pollution damage even without a decline in total emissions. These considerations motivate three questions. How do abatement and pollution shifts jointly determine emissions reductions? How do such regulations affect consumption and pollution damage across space? How would changes in geographic coverage and spatial variation in pollution prices alter these outcomes?

To answer these questions, I study South Korea's local air pollutant cap-and-trade system, the *Total Volume Control of Pollutants at Place of Business*. Expanded geographically in 2020, the policy regulates large polluting manufacturing and power plants emitting nitrogen oxides (NO_x), sulfur oxides (SO_x), and total suspended particulates (TSP) in polluted, densely populated districts. Between 2020 and 2024, emissions caps summed across the policy's four regulated areas declined by 28.6% for NO_x, 32.7% for SO_x, and 26.7% for TSP. These caps apply to covered establishments within regulated districts rather than all national emissions, so aggregate outcomes also depend on responses at unregulated plants within and outside regulated districts.

To analyze this policy, I combine nationwide administrative records and censuses for 2016–2023 to examine abatement and pollution shifts in manufacturing and power generation. District–industry emissions records capture process and combustion emissions and account for emission-control equipment, revealing changes that fuel use alone would miss. Linked plant-level records allow me to study abatement through energy and equipment

[1] Examples include bans on the use of certain fuels, plant-level emissions caps, limits on pollutant concentrations in exhaust gases, and cap-and-trade systems.

use and trace production reallocation through sales and input use. Nationwide coverage makes it possible to find a comparable control group and to follow these adjustments within and across firms, sectors, and districts.

Empirically, I use a difference-in-differences approach at the district–industry level to compare emissions changes across regulated and unregulated districts within industries and examine pollution shifts. Manufacturing emissions of NO_x, SO_x, and TSP fall by 35.5%, 36.9%, and 32.9%, respectively, in regulated districts relative to the same industries in unregulated districts. Under parallel counterfactual trends, these estimates combine the policy's effects on both groups: these relative declines can reflect emissions reductions in regulated districts, emissions increases in unregulated districts, or both. They do not identify the effect on either group separately or the national emissions change. Among unregulated districts, manufacturing emissions increase more in those closer to regulated districts. These estimates provide evidence of geographic pollution shifts: nearby unregulated plants can absorb displaced demand while facing lower transport costs to serve the affected markets than more distant plants.

Plant-level comparisons connect these emissions patterns to evidence on abatement and production reallocation. Relative to matched unregulated unaffiliated plants, regulated manufacturers exhibit larger increases in electricity use, electricity's share of energy use, and emission-control equipment use. These responses show why production changes alone do not capture the emissions consequences of regulation: abatement can reduce emissions intensity, while electrification can shift emissions to upstream power plants. Furthermore, unregulated affiliates exhibit stronger sales growth than unaffiliated plants in the same industries, providing evidence of within-firm production reallocation.

To account for abatement and pollution shifts and to translate *relative* responses into aggregate and spatial effects, I develop a spatial equilibrium model of manufacturing and upstream power generation. The spatial structure captures geographic variation in policy exposure and location-dependent pollution damages, while the equilibrium structure captures the market adjustments that reallocate production and emissions across plants, sectors, and regions. The model also allows me to evaluate alternative policies.

The model links process and combustion emissions to plants' input choices, equipment use, and production allocation. Pricing emissions changes both emissions intensity and the relative production costs that govern reallocation, making abatement and reallocation jointly determined. Electrification creates a further link: reducing direct manufacturing emissions can increase emissions from upstream power generation.

I discipline the model in two steps. First, I use pre-regulation data to determine the production and demand structure. Second, conditional on observed permit prices, I

estimate parameters governing abatement and effective compliance costs. This second stage targets the estimated responses to the policy together with cap-reduction benchmarks. The model closely reproduces the targeted emissions responses and captures the broad spatial distribution of emissions, which is not targeted in estimation.

Using this model, I obtain two main quantitative results. First, geographic pollution shifts account for 19.1% of the manufacturing emissions difference-in-differences estimate. The model provides this structural interpretation of the empirical evidence by decomposing the responses of regulated and unregulated district–industry cells. Second, the actual policy reduces aggregate pollution damage by 37.8% at the cost of a 1.42% decline in consumption. Damage reductions in regulated districts, where baseline pollution damage is concentrated, outweigh increases elsewhere. Consumption falls in both groups, with a slightly larger decline in regulated districts.

These results make geographic coverage and spatial variation in pollution prices central policy questions. I compare two nationwide alternatives that target the actual policy's national emissions totals for each pollutant. Nationwide coverage limits geographic pollution shifts and broadens the set of plants contributing to these targets. Under uniform pricing, the model reaches the targets with a smaller consumption loss than under the actual policy. However, smaller damage reductions in originally regulated districts, where baseline damage is concentrated, outweigh the additional benefits in previously unregulated districts. The second alternative retains nationwide coverage and sets prices proportional to local marginal damages (Muller and Mendelsohn 2009; Shapiro and Walker 2025). It directs stronger relative incentives toward locations with higher marginal damages, achieving greater aggregate damage reductions than either the actual policy or uniform pricing, with a smaller consumption loss than the actual policy.

Related literature This paper contributes to the literature on how environmental regulation affects emissions and economic activity. One part of this literature studies abatement through changes in production technologies, energy inputs, and emission-control investment (Fowlie 2010; Blundell et al. 2020; De Simone et al. 2024). A complementary literature studies pollution leakage and sectoral emissions shifts, including international reallocation associated with the pollution haven effect (Levinson and Taylor 2008; Martin et al. 2014; Naegele and Zaklan 2019), sectoral reallocation between manufacturing and power generation (Curtis and Lee 2019; von Graevenitz and Rottner 2026), and regional and intra-firm reallocation (Chen et al. 2025; Jia et al. 2025).[2] Direct empirical evidence on increases in production or emissions in unregulated locations induced by environmental

[2]I reserve the term leakage for policy-induced shifts in emissions from regulated to unregulated plants.

regulation remains limited. Examples include cross-border displacement in Tanaka et al. (2022), emissions increases outside regulated regions in Zhang and Zhao (2023) and Jiang et al. (2024), and production reallocation through firm networks and neighboring regions in Curtis et al. (2026).[3]

Closely related studies evaluate pollution markets: NO_x trading under Southern California's REgional CLean Air Incentives Market (RECLAIM) (Fowlie et al. 2012), China's regional carbon-market pilots (Cui et al. 2021), the European Union Emissions Trading System (Dechezleprêtre et al. 2023; Colmer et al. 2025), and PM trading in Surat, India (Greenstone et al. 2025). Together, these studies highlight two empirical requirements central to my analysis: a plausible counterfactual for regulated plants and firms and consistent data on emissions, abatement, and production.

I complement these evaluations by combining evidence on emissions, abatement, and production reallocation in Korea. Using nationwide emissions records, I provide evidence of emissions reductions and the corresponding geographic pollution shifts: manufacturing emissions increase more in unregulated districts closer to regulated districts. Plant-level records complement this spatial evidence by revealing abatement responses, within-firm production reallocation, and potential emissions displacement to power generation. Examining these outcomes together matters because abatement and production reallocation have different implications for pollution and production. I combine this evidence with a spatial equilibrium model to quantify these adjustments' contributions to emissions changes and their joint effects on pollution damage and consumption, and to compare alternative policy designs.

This structural approach relates to the literature that uses models to study the spatial aspects of environmental policy. Several studies develop theoretical frameworks (Weisbach et al. 2023; Kortum and Weisbach 2024), while others quantify the effects of policies on abatement and leakage across regulatory boundaries (Fowlie 2009; Fowlie et al. 2016; Xiang 2023; Farrokhi and Lashkaripour 2025; Cruz and Rossi-Hansberg 2025; Hollingsworth et al. 2026). Arkolakis et al. (2025) study long-run industrial decarbonization through small and medium-sized firms' lumpy decisions to adopt or drop fuels under regulations and subsidies. I complement their analysis by examining how short-run marginal abatement and production reallocation of large polluters jointly shape local air pollution under a specific place-based cap-and-trade program.

I build the model on two frameworks. Shapiro and Walker (2018) model manufacturing firms' abatement and reallocation responses to pollution taxes. Arkolakis and Walsh

[3]Related work examines the physical transport of pollution across regions and countries (Heo et al. 2025).

(2023) link manufacturing to upstream power generation and study how technological improvements and subsidies expand clean renewable generation capacity and achieve clean growth. Like Chen et al. (2025) and Curtis et al. (2026), I use an equilibrium model to interpret the reduced-form estimates as differences in responses between regulated and unregulated units and to quantify aggregate effects. The former models firm-level energy-use quotas, while the latter represents abatement through a single effort choice and summarizes regulatory costs through effective location productivity.

Relative to Shapiro and Walker (2018), Chen et al. (2025), and Curtis et al. (2026), I distinguish process and combustion emissions and model factor substitution, fuel substitution, and emission-control equipment as separate adjustment margins. The model further allows production and emissions to reallocate across plants within firms and across sectors. Relative to Arkolakis and Walsh (2023), I hold clean generation capacity fixed and study the short-run emissions consequences of regulation-induced industrial electrification.[4] In this setting, electrification can reduce manufacturers' direct emissions while shifting emissions to upstream power generation. These choices matter because abatement changes both emissions intensity and the production costs that govern reallocation, while electrification can shift emissions upstream and hence affect pollution damage and consumption across space.

The remainder of the paper is structured as follows. Section 2 describes the policy and data. Section 3 estimates relative emissions responses and examines adjustment mechanisms. Sections 4 and 5 develop and quantify the model. Section 6 evaluates the actual policy and nationwide alternatives, and Section 7 concludes.

2 Institutional Setting and Data

This section explains the policy background and describes the main dataset used in the analysis.

2.1 Policy Background

Local air pollutants, such as NO_x, SO_x, and TSP, harm human health by exacerbating asthma, bronchitis, and heart disease and increasing the risk of premature death (see, e.g., Gillingham and Huang 2024). They also damage buildings and infrastructure through

[4] I classify biofuel, coal gasification, fuel cell, hydro, marine (ocean), nuclear, solar, waste, and wind as sources of clean power generation and coal, gas, and oil (fossil fuels) as sources of dirty power generation. Apart from nuclear, these sources follow the Korean classification of new and renewable energy.

corrosion, contribute to smog and acid rain, reduce visibility, harm crops and forests, and degrade ecosystems. Thus, Korea has regulated them since the early 1990s with two main instruments: command-and-control regulation of pollution concentrations and pollution taxes. Neither instrument, however, directly constrained the total quantity of emissions. Hence, the government introduced cap-and-trade for local air pollutants.

It began in the Seoul Capital Area in 2007, initially covering NO_x and SO_x, and added TSP in 2018. In 2020, it expanded to the Central, Southern, and South-East Areas, leaving the remaining districts unregulated (Figure 1). I focus on this geographic expansion in 2020, which increased the number of regulated establishments across all covered sectors from roughly 500 to 1,600. Treated establishments were exempt from part of the pre-existing pollution tax to avoid double regulation because the cap-and-trade system is more stringent.[5] The control group consists of plants that were never brought into the new system.

The program reflects two government priorities. First, as shown in Figure 1, it targets polluted and densely populated districts since the effects of the regulated pollutants are concentrated in space (see Table A.1 for their spatial variation). When announcing the policy, the Ministry of Environment reported that the designated areas accounted for 88% of the national population and 38% of land area in 2018. The Ministry's modeling attributed 82% of the domestic contribution to $PM_{2.5}$ concentrations to emissions from these areas.[6] Several districts with substantial pollution damage remain unregulated, including districts in Gangwon Province located at the upper-right corner of the maps.

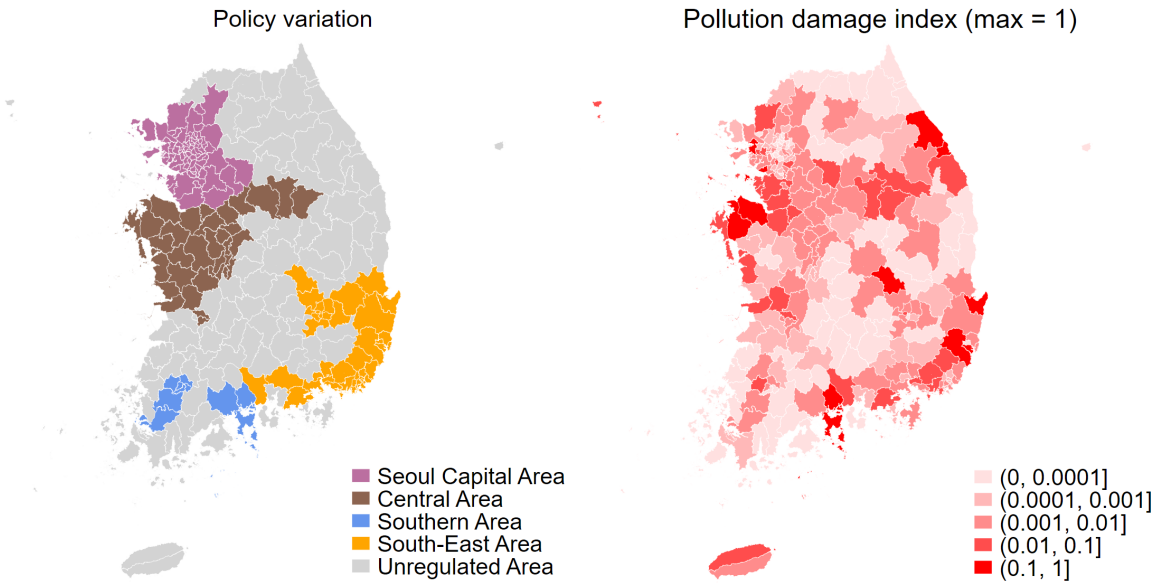
The boundaries of policy areas do not coincide with standard administrative boundaries, which helps distinguish the expansion from other regional policies. These boundaries define four policy areas and, combined with the three regulated pollutants, 12 separate permit markets. Caps are set separately for each market and decline over time; the left panel of Figure 2 shows their sums across pollutants within each area. Cap-reduction schedules differ across markets and are set jointly by the central and local governments.

Within these caps, the program freely allocates emissions permits to regulated establishments, with initial allocations based on pre-regulation emissions. Consistent with regional reduction goals, permits can be traded only within these markets. The right panel of Figure 2 plots transaction-volume-weighted average permit prices by area over time. Given the declining path of permit supply, initially high but falling average prices are con-

[5] In addition to the partial tax exemption, regulated establishments generating at least 10 but less than 20 tons of air pollutants annually before abatement remain subject to the pre-existing concentration requirements, but face less stringent limits.

[6] According to Statistics Korea, in 2018, 93% of manufacturing and power plant sales were concentrated in regulated districts.

Figure 1: Geographic coverage



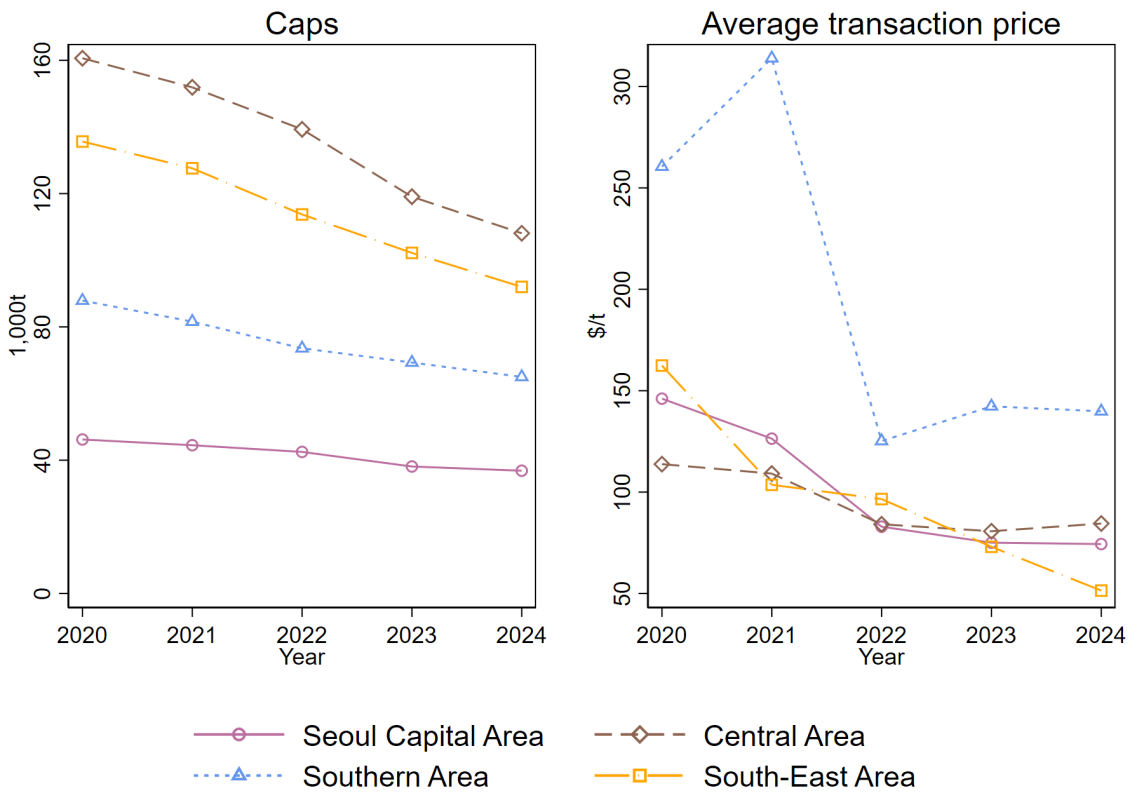
Notes: The left panel shows policy coverage across South Korea's 250 districts, which are the smallest self-governing administrative units within 17 metropolitan cities and provinces and serve as the primary administrative level for implementing local air quality policies. The right panel shows a baseline pollution-damage index constructed by multiplying each district's 2018 manufacturing and power generation emissions by pollutant-specific unit damage values, summing across pollutants, and weighting by district population. The index is normalized so that the largest district value equals one.
Sources: Ministry of Environment, Statistics Korea.

sistent with emissions reductions that lowered permit demand more rapidly than the cap. Prices later stabilized rather than collapsing, suggesting that the declines in permit supply and demand became more closely aligned while the market remained active.[7] Although the markets cover establishments outside manufacturing and power generation, permit prices reflect the marginal compliance incentives faced by all covered plants. Because every regulated plant can trade, these prices are informative even for plants that do not transact.

The program's second priority is to concentrate regulation on large stationary sources, which it does through emissions thresholds. According to the Ministry of Environment, business establishments contributed about 39% of ambient $PM_{2.5}$ pollution. Existing establishments are selected through a two-stage rule. First, their combined annual generation of

[7]Figures A.3–A.5 report trading volumes and transaction counts, overall and by pollutant.

Figure 2: Caps and permit prices



Notes: The left panel sums NO_x, SO_x, and TSP caps within each policy area. The right panel reports annual transaction prices averaged across these pollutants using paid-trade-volume weights. Prices exclude transactions without payment and are converted from Korean won to U.S. dollars using annual average exchange rates. There is no explicit price-stabilizing mechanism. Both panels cover 2020–2024; the earlier phase in the Seoul Capital Area is omitted. Table A.2 reports market-specific caps and reduction rates; Figures A.8–A.9 show cap and price paths by pollutant.
Sources: Ministry of Environment, Federal Reserve Bank of St. Louis.

NO_x, SO_x, and TSP before abatement must be at least 10 tons. Second, pollutant-specific regulation applies when annual emissions exceeded 4 tons for NO_x, 4 tons for SO_x, or 0.2 tons for TSP in either of the preceding two years. Although this rule is sector-neutral ex ante, manufacturing and power generation account for 82% of establishments in the policy registry. This reflects the fact that, among regulated sectors in regulated districts, manufacturing and power generation accounted for 73%, 93%, and 96% of NO_x, SO_x, and TSP emissions, respectively, in 2018.[8] Lastly, the program also requires plans for

[8]In general, regulated stationary sources—especially power generation and manufacturing and industrial facilities—are central sources of emissions: in India, thermal power plants account for about half of national SO_2 emissions and industrial activities another third, while fossil-fuel

emission-control equipment and its maintenance tailored to each emission source.

Geographic coverage and plant-level emissions thresholds jointly determine regulatory exposure and create several margins of adjustment. Since both manufacturing and power plants produce tradable goods, they can produce in one location and sell in others. Since electricity is a clean energy source from the perspective of manufacturing plants, they can use more electricity and less fossil fuel. Thus, abatement and pollution shifts can arise, leading to different distributional consequences across space. Indeed, interviews with government officials at the Ministry of Environment suggest that policymakers were aware of the possibility of leakage.[9]

A central question is therefore how much the policy induced abatement and pollution shifts within and across firms, sectors, and districts. For additional institutional detail, see Appendix A.1.

2.2 Data

I combine administrative records and censuses on emissions, economic activity, and abatement for manufacturing and power plants from 2016 to 2023.[10] The Korean Ministry of Environment records provide the registry of regulated plants, permit-market outcomes, and emissions data from the Tele-Monitoring System (TMS) and the Clean Air Policy Support System (CAPSS).[11] Statistics Korea's Statistical Business Register (SBR) and Mining and Manufacturing Survey (MMS) provide plant-level sales, inputs, locations, and five-digit Korean Standard Industrial Classification (KSIC) codes for manufacturing plants. For manufacturing plants, the SBR reports fuel use and greenhouse gas emissions by fuel type (coal, electricity, gas, oil, renewables, and other fuels) based on the MMS sample. For power plants, I additionally use the Economic Census and data from the Korea Electric Power Corporation (KEPCO) on fuel use, generating capacity, and electricity output for

combustion in industry and industrial processes add more than 30% of direct $PM_{2.5}$ emissions in 2019; in China, industrial sectors account for 72.8–86.1% of anthropogenic PM and 74.3–91.0% of SO_2 emissions between 2010 and 2018; globally, energy supply and industry account for 34% and 24% of anthropogenic greenhouse gas emissions, respectively, in 2019 (International Energy Agency 2021; Intergovernmental Panel on Climate Change 2022; Tang et al. 2023).

[9] I interviewed officials in the Ministry of Environment's Air Quality Management Division, which oversees the program, and its Metropolitan Air Quality Management Office, which has administered the Seoul Capital Area program since 2007.

[10] Unlike long-lived greenhouse gases, the local air pollutants studied here have short atmospheric lifetimes, so changes in emissions can quickly alter pollution exposure. This makes near-term changes in the spatial distribution of emissions relevant for policy evaluation.

[11] Regulatory monitoring directly measures 63%, 45%, and 22% of NO_x, SO_x, and TSP emissions, respectively. These records measure emissions at domestic sources, not ambient concentrations that can include pollution transported from China.

larger power plants. I also use administrative records on emission-control equipment in manufacturing and power plants from the Ministry of the Interior and Safety and the Ministry of Environment.

I merge these datasets using tax identifiers, locations, and KSIC codes. I link 1,093 regulated manufacturing and power plants—940 manufacturing plants and 153 power plants—to the administrative and census data. The resulting nationwide panel covers 250 administrative districts, 111 emission-source industrial classifications, and 90 three-digit KSIC industries, dropping industries that do not have any regulated plants. Plant-level energy and equipment records provide evidence on abatement choices, while sales linked to ownership and location allow me to trace production reallocation within and across firms, sectors, and districts.

In addition, the data provide a precise measure of local air pollutant emissions. Unlike emissions imputed solely from fossil-fuel combustion, this measure distinguishes production-process emissions from fuel-combustion emissions. Emissions from fossil-fuel use account for 83%, 49%, and 58% of NO_x, SO_x, and TSP emissions, respectively, over the sample period. However, emissions from production processes still account for a significant portion. For example, paper-producing plants emit large amounts of TSP. I therefore retain plants that report no fossil-fuel use in the analysis. The emissions measures also account for the operation of emission-control equipment.

To support environmental policy implementation, the Stack Emission Management System (SEMS) collects plant- and stack-level information used to construct the CAPSS emissions inventory. I use publicly available CAPSS emissions aggregated to district–industry cells. At this level of aggregation, emissions in regulated district–industry cells include those of unregulated plants. Equipment use, inputs, and sales are observed at the plant level. I therefore study emissions at the district–industry level and abatement and reallocation mechanisms at the plant level. Appendix B describes data construction, record linkage, variable definitions, and sample restrictions.

3 Empirical Analysis

This section examines the effects of the policy on emissions and the underlying adjustment mechanisms. Under pollution shifts, the difference-in-differences estimates combine the policy's effects on regulated units with its spillover effects on unregulated units. For emissions, the estimated relative decline can therefore reflect both reductions in regulated districts and increases in unregulated districts; it is not the aggregate emissions reduction. To the extent that such shifts raise emissions in unregulated districts, these estimates can

overstate reductions in regulated districts.

For estimation, I employ a difference-in-differences approach, comparing trends between regulated and unregulated district–industry cells (d, s). Following Heckman et al. (1997) and Heckman et al. (1998), I also use matching to compare trends between regulated and unregulated plants (h). Pollutant p emissions, sales, employment, fuel use, and electricity shares are expressed in logarithms. Electricity share is electricity use divided by total electricity, coal, oil, and gas use, all measured in tons of oil equivalent. Equipment outcomes are an indicator for updating emission-control equipment and the product of this indicator and log annual equipment operating hours. The baseline event-study specification normalizes 2019 as the base period. The baseline empirical research design treats multi-pollutant regulations as a single composite policy. For more details, see Appendix C.

3.1 District–Industry Level

This section examines relative changes in pollutant p emissions across district d–industry s cells. I use observations with positive, non-missing emissions of all three pollutants for comparability.[12]

3.1.1 Specification

At the district–industry level, I compare the trends of emissions in regulated district–industry cells and those in unregulated district–industry cells, using:

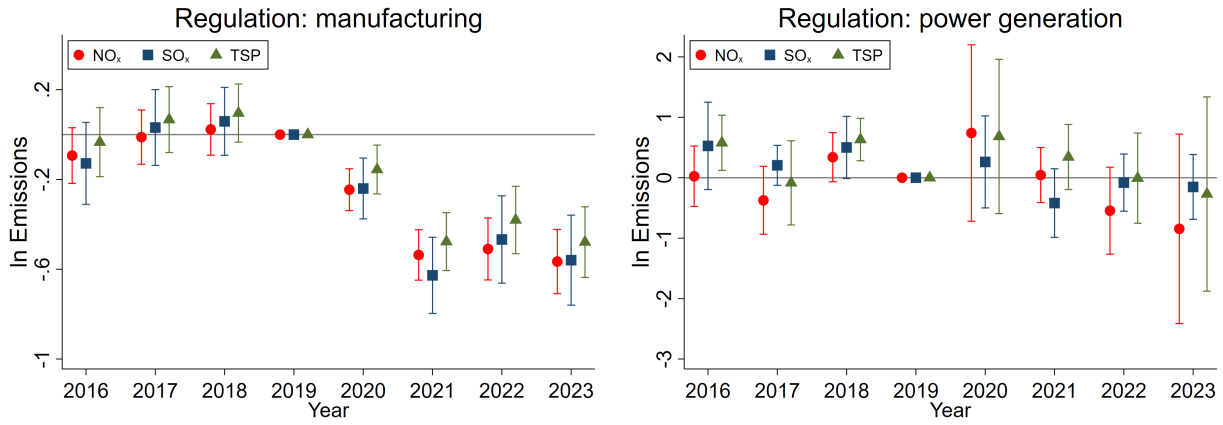
$$y^{p}_{d,s,t} = \alpha_{d,s} + \alpha_{s,t} + \alpha_{r,t} + \sum_{\tau \neq 2019} \beta_{\tau} \times Treat_{d} \times Year_{\tau} + \varepsilon_{d,s,t}, \tag{1}$$

$$y^{p}_{d,s,t} = \alpha_{d,s} + \alpha_{s,t} + \alpha_{r,t} + \beta \times Treat_{d} \times Post_{t} + \varepsilon_{d,s,t}, \tag{2}$$

for the event-study and pooled specifications, respectively. For power generation, industry–year fixed effects are omitted due to its small sample size. $Treat_d$ is a binary indicator equal to 1 if district d is regulated and 0 otherwise. Here, r denotes provinces, t indexes calendar years, $Year_\tau$ indicates calendar year τ, and $Post_t$ equals 1 for years after 2019. Standard errors are clustered at the district–industry level. I do not cluster at the policy area level

[12]This choice preserves the paper's log-emissions framework and estimates proportional changes among active source categories. I do not impose a strongly balanced panel as the main specification because, unlike plant-level balancedness, it would keep only district–industry cells that report positive emissions in every year rather than holding fixed well-defined plants. The corresponding balanced-panel estimates are reported in Appendix D.

Figure 3: Event dynamics: district–industry changes



Notes: These two panels show the event dynamics for manufacturing and power generation emissions across three regulated local air pollutants. The results are calculated from Equation (1). The 95% confidence intervals are shown jointly with the point estimates.
Sources: Ministry of Environment.

because there are only five areas including the unregulated one. Appendix D reports robustness checks with district-clustered standard errors. The coefficients of interest are β_τ and β, which summarize differences in emissions changes between regulated and unregulated districts for calendar year τ relative to 2019 and for the pooled post-policy period, respectively. Since only regulated industries are kept in the sample, the control group at the district–industry level is a set of industries in unregulated districts.

3.1.2 Results

Figure 3 and Table 1 summarize emissions responses at the district–industry level. First, manufacturing emissions of NO_x, SO_x, and TSP in regulated districts fall by 35.5%, 36.9%, and 32.9%, respectively, relative to the same industries in unregulated districts. For context, aggregate caps for covered establishments declined by 26.7–32.7% across pollutants between 2020 and 2024. The empirical estimates instead measure relative emissions changes during 2020–2023 across regulated and unregulated districts, including responses at unregulated plants, rather than reductions among covered establishments alone. The large relative declines are notable because emissions in regulated districts also include unregulated plants that may receive displaced production. The large declines in district–industry emissions are also consistent with regulated plants' substantial share of baseline emissions.

Second, the emissions response differs across sectors. The reductions are driven primarily by manufacturing, which accounts for a large share of baseline emissions (Table A.3)

Table 1: Policy effect: district–industry changes

	Manufacturing	*Power generation*
ln NO_x emissions	−0.44***	−0.15
	(0.05)	(0.17)
ln SO_x emissions	−0.46***	−0.41
	(0.08)	(0.25)
ln TSP emissions	−0.40***	−0.09
	(0.06)	(0.21)
N	39,086	1,219

Notes: The average effects are calculated from Equation (2). All specifications include district–industry and province–year fixed effects, plus industry–year fixed effects for manufacturing. Standard errors, clustered at the district–industry level, are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment.

and exhibits large, statistically robust declines in NO_x, SO_x, and TSP. Power generation estimates are imprecise, with event-study point estimates that vary in sign and become negative in later years.[13]

Electrification links manufacturing abatement to emissions from power generation. Regulated manufacturing plants can reduce on-site emissions partly by increasing electricity use, while power plants may meet the resulting demand. In Korea's short-run dispatch setting, lower-marginal-cost generators, including available clean generation, are already dispatched, so additional demand can be met by operating existing fossil-fuel plants more intensively without changes in generation capacity.

3.2 Geographic Pollution Shifts

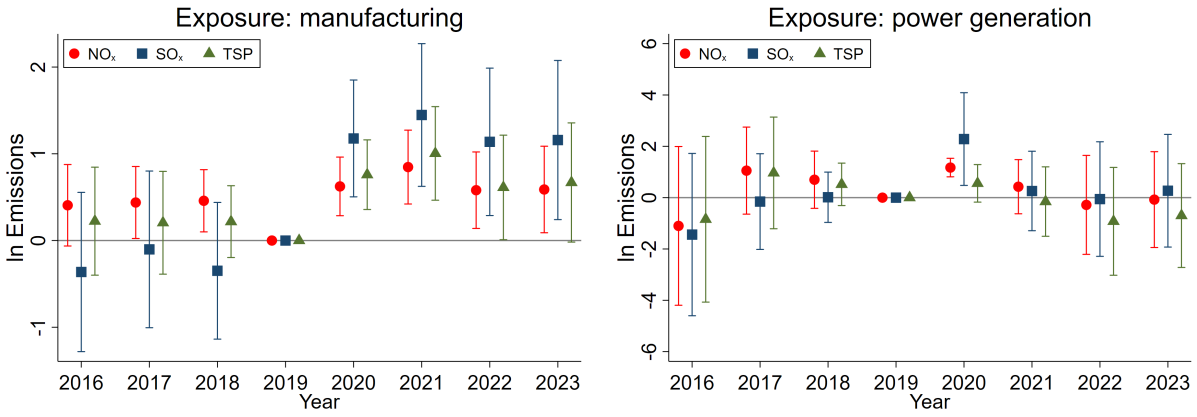
I next examine geographic pollution shifts by comparing emissions responses across unregulated districts with different exposure to nearby regulation.

3.2.1 Specification

Lower transport costs make nearby unregulated plants more likely destinations for displaced production. I therefore examine how emissions responses vary with distance from regulated districts. Let D_d denote the centroid-to-centroid distance, in kilometers, from

[13]Motivated in part by the small sample size in power generation, I report MacKinnon–Webb-style placebo randomization-inference p-values as a supplementary check in Table A.6 (MacKinnon and Webb 2020). The power generation estimates generally fail to reject the null under the placebo distribution, with only weak 10% evidence for SO_x; the manufacturing estimates remain highly significant under this placebo benchmark.

Figure 4: Event dynamics: geographic pollution shifts



Notes: The coefficients measure how emissions changes in unregulated districts vary with proximity to regulated districts. The results are calculated from Equation (3). The 95% confidence intervals are shown jointly with the point estimates.
Sources: Ministry of Environment.

unregulated district d to its nearest regulated district. I exclude districts with $D_d \leq 45$—roughly one average district diameter—because centroid-to-centroid distance provides a coarse measure of exposure at shorter distances. I define $\text{Exposure}_d \equiv -\ln(D_d)$ and compare trends across the remaining unregulated district–industry cells:

$$y^p_{d,s,t} = \alpha_{d,s} + \alpha_{s,t} + \sum_{\tau \neq 2019} \beta_\tau \times \text{Exposure}_d \times Year_\tau + \varepsilon_{d,s,t}, \tag{3}$$

$$y^p_{d,s,t} = \alpha_{d,s} + \alpha_{s,t} + \beta \times \text{Exposure}_d \times Post_t + \varepsilon_{d,s,t}. \tag{4}$$

For power generation, I replace industry–year fixed effects with year fixed effects due to its small sample size. Exposure is fixed over time, and higher values indicate greater proximity to regulated districts. The event-study coefficients measure differential changes in ln Emissions per unit increase in exposure relative to 2019. The pooled specification estimates the corresponding differential change before and after the policy. Standard errors are clustered at the district–industry level.

3.2.2 Results

Figure 4 and Table 2 summarize geographic pollution shift responses at the district–industry level. Manufacturing emissions rise more in unregulated districts closer to regulated districts than in those farther away. These estimates provide evidence of geographic pollution shifts, which can magnify the estimated relative emissions decline. For

Table 2: Policy effect: geographic pollution shifts

	Manufacturing	*Power generation*
ln NO_x emissions	0.35	0.07
	(0.19)	(0.67)
ln SO_x emissions	1.42***	1.05
	(0.41)	(0.74)
ln TSP emissions	0.61*	−0.52
	(0.26)	(0.95)
N	2,105	101

Notes: The average effects are calculated from Equation (4). The sample consists of unregulated districts more than 45 km from the nearest regulated district. Specifications include district–industry fixed effects and industry–year (manufacturing) or year (power generation) fixed effects. Standard errors, clustered at the district–industry level, are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment.

power generation, the geographic pollution shift estimates are imprecise, as in the baseline emissions analysis.[14] Lastly, note that this evidence of geographic pollution shifts, on top of the relative emissions declines found above, helps to address potential confounding from other environmental regulations because the overlapping policies discussed in Appendix D do not share the expansion's geographic targeting.

3.3 Plant Level

This section investigates the mechanisms behind such changes in emissions. I restrict the analysis to a balanced intensive-margin sample with non-missing sales, electricity, employment, and equipment-use information throughout the panel. This choice reflects the fact that the period of interest is 8 years and large polluters rarely exit. For power plants, electricity use and employment are omitted from the requirement because the former is not observed and the latter is often noisy. This setting avoids conflating changes in outcomes with changes in reporting or entry and exit. Lastly, I exclude from the sample plants belonging to firms whose first regulated plant appears only after the initial 2020 rollout. This restriction removes a few later-treated firms and ensures that treatment status is defined relative to the initial policy expansion.

[14] In Appendix D, a distance-ring specification also yields positive manufacturing emissions responses relative to the farthest ring, without imposing a linear relationship with Exposure_d.

3.3.1 Specification

At the plant level, I use:

$$y_{h,t} = \alpha_h + \alpha_t + \sum_{\tau \neq 2019} \beta_\tau \times Treat_h \times Year_\tau + \varepsilon_{h,t}, \tag{5}$$

$$y_{h,t} = \alpha_h + \alpha_t + \beta \times Treat_h \times Post_t + \varepsilon_{h,t}. \tag{6}$$

I first compare regulated with unregulated unaffiliated manufacturing plants. To examine pollution leakage within firms, I then compare unregulated affiliated with unregulated unaffiliated plants. In the latter, treated plants are unregulated affiliates of firms with regulated plants in the same three-digit KSIC industry. Industries are defined using pre-2020 KSIC codes to avoid endogenous reclassification. The control group therefore consists of unregulated unaffiliated plants in the same three-digit industry, matched using pre-policy size measures. Following the literature (Chen et al. 2025; Colmer et al. 2025; Curtis et al. 2026), I use relative production responses at unregulated manufacturing plants as a proxy for pollution leakage through production reallocation. This interpretation requires that differential changes in output prices and emissions intensity do not offset the sales response.

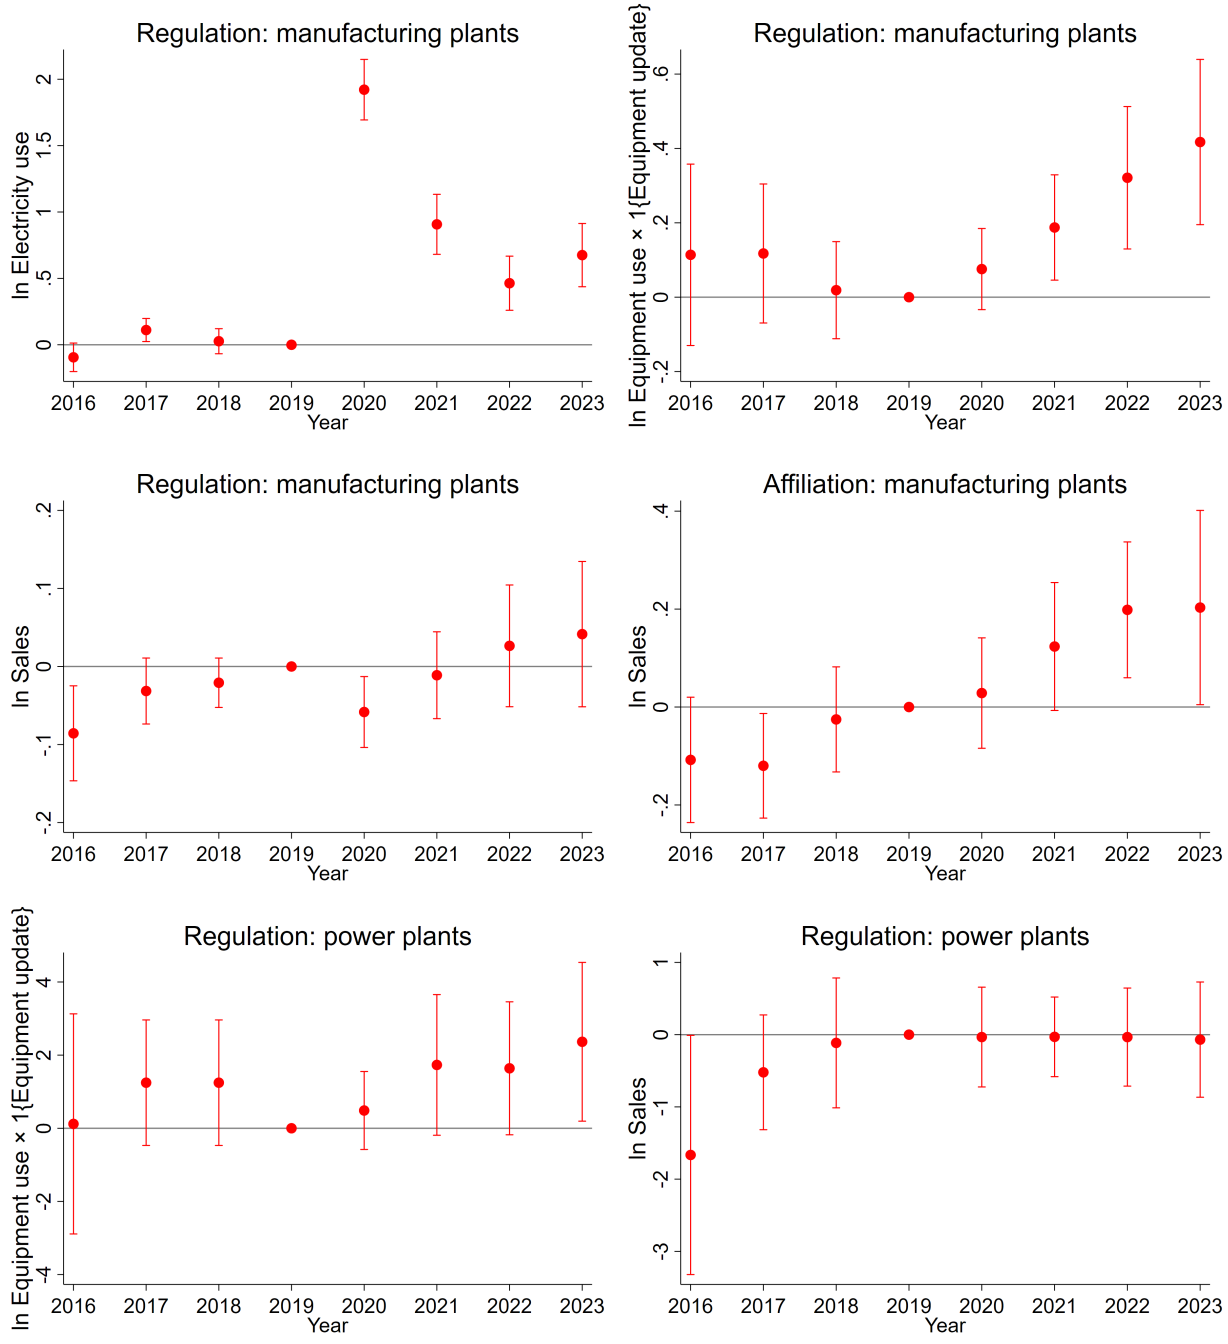
The distinction between affiliated and unaffiliated plants matters because large polluters often belong to large firms and have many unregulated affiliated plants under the same firm ownership. Of regulated manufacturing plants, 67% have an unregulated affiliate, and 63% have one in the same three-digit KSIC industry. Among these same-industry affiliates, 86% are located in regulated districts and are typically near the relevant threshold. Strong within-firm reallocation could therefore attenuate the effect of the policy within regulated districts. Within-firm reallocation is less direct in power generation because centralized dispatch allocates output across generators according to generation costs, rather than firms reallocating production across affiliates (see Appendix A.2). Hence, I compare regulated with unregulated power plants.

I use one-to-one nearest-neighbor matching without replacement for balance. Following Abadie and Spiess (2022), standard errors are clustered by matched pair. For affiliated-plant specifications, I additionally cluster by firm to allow correlated shocks across a firm's plants. See Appendix C.1 for more details.

3.3.2 Results

Figure 5 and Table 3 summarize plant-level evidence on abatement and pollution shifts. As shown in the top two panels, regulated manufacturing plants show two abatement mecha-

Figure 5: Event dynamics: plant responses



Notes: These panels show the event dynamics from Equation (5) for electricity use, update-weighted log equipment use, and sales at manufacturing and power plants. "Regulation" panels compare regulated plants with matched unregulated unaffiliated plants (unregulated plants for power generation); the "Affiliation" panel compares unregulated affiliates of regulated firms with matched unregulated unaffiliated plants. The 95% confidence intervals are shown jointly with the point estimates.
Sources: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea.

Table 3: Policy effect: plant responses

	Manufacturing		*Power generation*
	Regulated	Unregulated affiliated	Regulated
1{Equipment update}	0.02*	0.04	0.13
	(0.01)	(0.03)	(0.11)
ln Equipment use × 1{Equipment update}	0.17*	0.40	0.90
	(0.09)	(0.20)	(0.89)
ln Sales	0.05	0.19**	0.53
	(0.03)	(0.06)	(0.38)
ln Electricity use	0.95***	0.39*	
	(0.08)	(0.18)	
ln Electricity share	0.10**	0.09	
	(0.04)	(0.06)	
N	6,544	1,536	224

Notes: The average effects are calculated from Equation (6). Manufacturing columns compare regulated plants and unregulated affiliates of regulated firms, respectively, with matched unregulated unaffiliated plants; the power generation column compares regulated with matched unregulated power plants. Standard errors, clustered by matched pair (and additionally by firm in the affiliated column), are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001. *Sources*: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea.

nisms. First, electricity use at regulated manufacturing plants increases by 158.7% relative to matched unregulated unaffiliated plants, indicating persistent industrial electrification. This response can shift emissions to regulated or unregulated power plants. Shifts to regulated generators remain within the cap-and-trade system rather than escaping regulation. Shifts that raise emissions at unregulated generators constitute leakage.[15] Second, relative to matched unregulated unaffiliated plants, regulated plants are 2.3 percentage points more likely to update their equipment, and their update-weighted log equipment use increases by 0.17. Part of the increase in electricity use reflects the fact that emission-control equipment is operated with electricity.

Ministry of Environment reports document a sharp increase in large-scale abatement projects initiated by regulated plants after the policy took effect.[16] These investments predominantly target multi-pollutant control technologies, such as filters and scrubbers, rather than long-term R&D. These technologies rely heavily on electricity without directly

[15]The supplementary event study shows a relative increase in regulated manufacturers' electricity shares, which is consistent with substitution toward a cleaner energy source from manufacturers' perspective (the upper-left panel of Figure A.7). Changes in electricity use and shares do not establish the corresponding change in fossil-fuel use for the full sample, which includes plants reporting no fossil-fuel use.

[16]This information is available on the Ministry of Environment's website: Environmental Information Disclosure System.

generating additional on-site emissions. An official from the Ministry's Air Quality Management Division also reported that regulated plants substantially increased investment in emission-control equipment following the policy.

I interpret changes in fuel composition and equipment updating as evidence of abatement, although leakage may affect plants' production and equipment operation. The middle-left panel shows an initial relative decline in regulated plants' sales, followed by a narrowing of the gap.[17] In the middle-right panel, the sales of unregulated affiliated plants grow by 21.4% relative to unregulated unaffiliated plants. I interpret these responses as an initial reallocation toward unaffiliated competitors, followed by stronger reallocation toward unregulated affiliates within firms.[18]

The bottom panels show a positive but imprecisely estimated equipment-use response at regulated power plants and little evidence of a sales decline. Manufacturing electrification creates an additional source of demand for power plants, which can offset the contraction from their own regulation.

3.4 Taking Stock

To summarize, manufacturing plants drive emissions reductions: manufacturing emissions fall by 32.9–36.9% across pollutants in regulated districts, relative to the same industries in unregulated districts. The evidence indicates reallocation within manufacturing firms and across districts, while increased manufacturing electricity use can shift emissions to power generation. For robustness checks, see Appendix D.

These effects align with the permit-market patterns shown in Figure 2. Note the timing gaps: firms adjust their permit positions throughout the accounting year in response to production, abatement, and emissions outcomes from the previous year. Early declines in permit prices are consistent with permit demand falling faster than caps, alongside increased equipment use and electrification at regulated manufacturers. Permit prices later stabilized, while the event study shows a stronger relative sales response at unregulated affiliated plants. The plant-level estimates show earlier abatement responses and a later strengthening of within-firm reallocation. The permit-market series alone do not distinguish these mechanisms.

The findings above resonate with the broader literature. The geographic pollution shift estimates complement evidence of production and emissions shifting to unregulated

[17] The accompanying decline in regulated manufacturers' employment (the upper-right panel of Figure A.7) provides additional evidence of a contraction in real activity.

[18] The lower panels of Figure A.7 show additional responses in unregulated affiliated manufacturing plants.

locations (Tanaka et al. 2022; Zhang and Zhao 2023; Jiang et al. 2024). Stronger leakage within firms than across firms is a pattern found in the literature (Chen et al. 2025; Curtis et al. 2026). Emissions shifts from manufacturing to power generation connect to the emerging literature on electrification and its environmental consequences (Curtis and Lee 2019). A key policy implication is that, in the short run, industrial electrification can induce unintended emissions from fossil-fuel power plants. The coexistence of abatement and pollution shifts differs slightly from the theoretical predictions of Xiang (2023) and Arkolakis et al. (2025), where pollution shifts dominate in the short run and clean innovation and large-scale abatement drive long-run decarbonization. This is because marginal abatement and technology adoption can occur more quickly for local air pollutants, relative to clean innovation and large capital investment.

These estimates capture differences in responses to the policy, including spillovers to unregulated units. On their own, they do not quantify aggregate benefits and costs or the effects of alternative policy designs. For example, pollution shifts to low-damage districts can decrease aggregate pollution damage because the effect of local air pollution is concentrated in space. I next use a spatial equilibrium model to address these limitations.

4 Quantitative Spatial Model of Manufacturing and Power Generation

I develop a spatial equilibrium model because policy exposure and pollution damages vary across regions and pollution shifts change emissions and production beyond regulated plants. I integrate the manufacturing and emissions model of Shapiro and Walker (2018) with the manufacturing and power generation model of Arkolakis and Walsh (2023). I extend these frameworks to distinguish process and combustion emissions and allow manufacturing plants to abate through factor substitution, fuel substitution, and emission-control equipment. These abatement choices interact with production reallocation within and across firms and districts. Power plants cannot substitute factors or fuels in the short run. Electrification links manufacturers' abatement to power-sector emissions.

The model is static because the data and policy horizon are short, local air pollutants are short-lived, and existing abatement technologies can be adopted more readily than greenhouse gas abatement technologies.[19] In a similar vein, I calibrate firm masses to

[19] A dynamic model can capture long-run clean innovation (Xiang 2023) or lumpy technology switching (Arkolakis et al. 2025). I focus on adjustments in fuel use and the operating intensity of emission-control equipment—end-of-pipe controls, such as filters and scrubbers—and abstract from innovation, large-scale abatement investment, equipment installation, and non-electric operating costs. These intensive margins are relevant to pollution taxes and emissions trading, which change large polluters' marginal incentives.

pre-regulation data and hold the establishment network fixed to isolate abatement and production reallocation within and across existing firms.[20] See Appendix E for derivations.

4.1 Environment

The economy consists of N discrete districts, each with a representative household, and two sectors: manufacturing and power generation. Manufacturing is divided into J sectors based on product similarity, and power generation into F sectors based on energy source. Throughout the analysis, subscripts denote districts, manufacturing sectors, and power generation fuel types, while superscripts index factor types, fuel types, and pollutants.

4.2 Manufacturing Firms

In district i and sector j, there is a fixed mass of manufacturing firms, $M_{i,j}$. A firm in this group draws productivity ϕ and supplies a differentiated good under monopolistic competition (Melitz 2003). This good is an aggregate of plant varieties within the firm. The firm can produce in places l other than its headquarters district and sell to destination districts n. This multiregional production feature addresses within-firm leakage. Each plant is indexed by its headquarters location i, production location l, sector j, firm productivity ϕ, and plant productivity z.

4.2.1 Manufacturing Production and Emissions

For simplicity, I drop the sector subscript j in the remaining manufacturing section. For a fixed output requirement of $\bar{q} = 1$, the plant chooses inputs and equipment intensity to minimize production cost:

$$\mu_{il}(\phi, z) = \min_{k,e,\ell;\{e^f\}_{f\in\mathbb{F}},c\in[0,1]} \left\{ rk + \sum_{f\in\mathbb{F}} p^f e^f + w_l \ell + p^e \delta c \bar{q} + \sum_{p\in\Lambda_{il}(\phi,z)} \tau_l^p m_{il}(\phi, z)\, s_{il}^p \right\}$$

$$\text{s.t. } \bar{q} = \phi z k^{\alpha^K} e^{\alpha^E} \ell^{\alpha^L},$$

$$e = \left(\sum_{f\in\mathbb{F}} \alpha^f \left(e^f\right)^{\frac{\lambda-1}{\lambda}} \right)^{\frac{\lambda}{\lambda-1}},$$

$$s_{il}^p = (1+c)^{-\alpha^p} \left(\kappa^p \bar{q} + \sum_{f\in\mathbb{F}} \kappa^{p,f} e^f \right).$$

[20] Figure A.2 shows no significant change in regulated districts' plant shares.

Here $\alpha^K, \alpha^E, \alpha^L > 0$, $\alpha^f > 0$, $\lambda > 1$, $p \in \{\text{NO}_x, \text{SO}_x, \text{TSP}\}$, $\alpha^p > 0$, and $\delta > 0$. Specifically, the plant combines capital k, energy e, and labor ℓ using a Cobb–Douglas production function. Energy is a constant-elasticity-of-substitution (CES) composite over fuel types $f \in \mathbb{F} = \{\text{Coal, Electricity, Gas, Oil}\}$ measured in tons of oil equivalent (toe). Electricity is the key fuel because it is used for both production and abatement in the model and universally observed across plants and years in the data. The parameters α^K, α^E, and α^L are output elasticities in the outer Cobb–Douglas nest, whereas α^f is a weight in the inner CES fuel aggregate. r is the rental rate of capital, p^f is the price of fuel f, and w_l is the wage. The parameter $\delta > 0$ measures electricity required per unit of output when emission-control equipment operates at full intensity. A plant producing $\bar{q}$ units of output at equipment intensity c uses $\delta c \bar{q}$ units of electricity to operate the equipment, in addition to production electricity e^e. Its total electricity use is therefore $e^e + \delta c \bar{q}$.

The Cobb–Douglas exponents sum to one, $\alpha^K + \alpha^E + \alpha^L = 1$, ensuring constant returns to scale (CRS). Energy is an essential input, since output is zero without it; Ryan (2018) finds even stronger complementarity between energy and other inputs. I normalize the CES fuel weights such that $\sum_{f \in \mathbb{F}} \alpha^f = 1$; this fixes the scale of the energy composite. At a given equipment intensity, production inputs, equipment electricity use, and emissions scale proportionally with output. Minimizing over equipment intensity preserves this proportionality, so unit cost equals marginal cost. I therefore evaluate the problem at $\bar{q} = 1$ to obtain the cost of producing one unit.

Emissions of pollutant p, s_{il}^p, consist of production process emissions and fossil-fuel use emissions. κ^p denotes the emission factor of pollutant p associated with the production process of this sector, while $\kappa^{p,f}$ represents the emission factor of pollutant p for fuel f. For example, if a plant produces paper and emits particulates into the air, κ^p captures it. If a plant uses fossil fuels heavily, $\kappa^{p,f}$ captures the corresponding emissions.

Each plant faces up to three local air pollutant regulations, based on its emissions volume and district. $\hat{s}_{il}^p(\phi, z)$ denotes pre-regulation emissions and $\bar{T}^p$ the threshold for pollutant p. The regulation is applied only to plants whose unregulated emissions meet or exceed the thresholds: $\hat{s}_{il}^p(\phi, z) \geq \bar{T}^p$. Permits are allocated in lump sum, based on the pre-regulation emissions discounted by the cap changes. Let $A_{il}^p(\phi, z)$ denote the initial allocation of permits for pollutant p to the plant. With tradable permits, this allocation affects plant profits but not marginal input or abatement choices. Permit prices τ_l^p vary by pollutant and district. The set of regulated pollutants, $\Lambda_{il}(\phi, z)$, depends on the plant's location and baseline emissions, which in turn depend on its headquarters location and productivity.

Implementing additional emission controls under cap-and-trade can require adjust-

ments to equipment and operating procedures that differ across production processes. The multiplier $m_{il}(\phi, z)$ approximates differences in incremental compliance burdens not captured by the modeled equipment effectiveness and electricity requirements. I use pre-policy emission-control equipment operating hours to proxy differences in compliance conditions across emissions sources and production processes. I assume that more established operation of emission controls lowers the incremental compliance burden. The multiplier is common across pollutants and emissions sources within a plant and leaves physical emissions accounting unchanged.

Once regulated, plants can reduce emissions by changing their fuel mix and operating emission-control equipment. Electricity produces no emissions on site and powers emission-control equipment, so regulation can induce electrification. Equipment intensity, $c_{il}(\phi, z) \in [0, 1]$, reduces emissions at an additional variable electricity cost, as with filters and scrubbers. Because equipment electricity use and pre-abatement emissions scale proportionally with output, optimal equipment intensity is independent of current output scale conditional on productivity, input prices, and regulatory exposure, but depends on emissions per unit of output, in the spirit of Copeland and Taylor (2004), Shapiro and Walker (2018), and Farrokhi and Lashkaripour (2025). Distinguishing process from combustion emissions allows regulation to affect plants through different abatement margins. Motivated by the differences in input use and abatement equipment documented in Table A.5, I model higher productivity as lowering fuel use and combustion emissions per unit of output while raising total emissions. Within a sector, this lowers the marginal benefit of using equipment.

Consistent with the operating details verified in the data, I represent the plant's emission-control equipment by a single operating intensity c. These end-of-pipe controls treat exhaust gases from production processes and fuel combustion before release. I allow higher c to reduce all three pollutants, at the cost of additional electricity use. The parameter α^p governs the effectiveness of emission-control equipment, while λ governs substitution across fuels. These parameters shape regulated plants' abatement choices and the resulting shifts in emissions. For example, λ also governs the potential shift from manufacturing to power generation.

The cost minimization problem can be solved as follows. At a given equipment intensity, substituting the emissions equation into the cost-minimization objective adds combustion-related compliance costs to each fossil fuel's purchase price, yielding the effective fuel price:

$$\tilde{p}_{il}^{f}(\phi, z) = p^{f} + \sum_{p \in \Lambda_{il}(\phi, z)} \tau_{l}^{p} m_{il}(\phi, z) \left(1 + c_{il}(\phi, z)\right)^{-\alpha^{p}} \kappa^{p,f},$$

for $f \neq e$ and, since electricity does not generate on-site emissions, its effective fuel price is:

$$\tilde{p}_{il}^{e}(\phi, z) = p^{e}.$$

Minimizing fuel expenditure for a given quantity of the energy composite yields its unit cost:

$$p_{il}^{E}(\phi, z) = \left(\sum_{f \in \mathbb{F}} \alpha^{f} \left(\frac{\tilde{p}_{il}^{f}(\phi, z)}{\alpha^{f}} \right)^{1-\lambda} \right)^{\frac{1}{1-\lambda}}.$$

Conditional on the use of emission-control equipment $c_{il}(\phi, z)$, the regulation affects input choices through fuel-specific effective prices. These effective prices determine the fuel composition, energy price index, and substitution between energy, capital, and labor. Substituting the cost-minimizing factor demands into the objective gives the unit production cost:

$$\underbrace{\mu_{il}(\phi, z)}_{\text{Unit production cost}} = \underbrace{\frac{G_{il}(\phi, z)}{\phi z}}_{\text{Factor cost}} + \underbrace{p^{e} \delta c_{il}(\phi, z)}_{\text{Equipment electricity cost}} + \underbrace{\sum_{p \in \Lambda_{il}(\phi, z)} \tau_{l}^{p} m_{il}(\phi, z) \left(1 + c_{il}(\phi, z)\right)^{-\alpha^{p}} \kappa^{p}}_{\text{Pollution cost from production process}}, \tag{7}$$

where the Cobb–Douglas factor cost index is:

$$G_{il}(\phi, z) = \left(\frac{r}{\alpha^{K}} \right)^{\alpha^{K}} \left(\frac{p_{il}^{E}(\phi, z)}{\alpha^{E}} \right)^{\alpha^{E}} \left(\frac{w_{l}}{\alpha^{L}} \right)^{\alpha^{L}}.$$

Appendix E provides the complete derivation.

I assume centralized pricing by atomistic firms. Each firm jointly chooses prices for all of its plant varieties to maximize total profits, taking sectoral price indices and expenditures as given and internalizing substitution among its plants. Given Dixit–Stiglitz preferences, the price of a plant variety is $p_{iln}(\phi, z) = \mathcal{M} \gamma_{il} \mu_{il}(\phi, z) t_{ln}$, where $\mathcal{M} = \sigma / (\sigma - 1)$ is the sector-specific markup, σ is the elasticity of substitution across firm composites, and $\eta > \sigma > 1$ is the elasticity of substitution across plant varieties within a firm. The price contains two spatial cost shifters: γ_{il} denotes the multiregional production frictions (e.g., communication or monitoring frictions), and t_{ln} denotes the iceberg trade costs (e.g., transportation costs). Let $R_{iln}(\phi, z) = p_{iln}(\phi, z) q_{iln}(\phi, z)$. Given this pricing rule and the preferences, the plant's profit is the sum of destination-specific profits, and the firm's

global profit is the sum of profits across its plants:

$$\pi_{il}(\phi, z) = \sum_{n} \frac{R_{iln}(\phi, z)}{\sigma} + \sum_{p \in \Lambda_{il}(\phi,z)} \tau_l^p A_{il}^p(\phi, z),$$
$$\pi_i(\phi) = \sum_{l} \mathbb{E}_z[\pi_{il}(\phi, z)].$$

Mechanisms and Empirical Implications I characterize manufacturing adjustment in a conditional equilibrium that holds factor, fuel, and electricity prices, and final expenditures fixed. I absorb each plant's compliance-cost multiplier into its effective pollution prices. This analysis addresses two questions raised by the empirical results. First, how does a higher pollution price reduce emissions at a regulated manufacturing plant? Second, why can the same policy raise economic activity and emissions at unregulated plants?

I fix a manufacturing plant in district i, suppress the other indices, and consider a local increase in τ_i^p, holding all other permit prices fixed. The price change applies to all plants covered by that pollutant's regulation in the district. Note that, because the common emission-control equipment c_i affects emissions of every active pollutant, the effective pollution price wedges for all active pollutants adjust endogenously. Then, the local response of plant-level emissions to τ_i^p admits the following accounting decomposition:

$$\underbrace{d \ln s_i^p}_{\text{Changes in emissions}} = \underbrace{\Gamma_i^p d \ln \tau_i^p}_{\text{Emission control}} + \underbrace{\sum_{f \neq e} \psi_i^{p,f} \mathcal{W}_i^{p,f} d \ln \tau_i^p}_{\text{Fuel substitution}} + \underbrace{\Xi_i^p d \ln \tau_i^p}_{\text{Electricity intensity}} + \underbrace{\Upsilon_i^p d \ln \tau_i^p}_{\text{Scale}}.$$

The first three terms are within-plant margins: Γ_i^p is the contribution of emission-control equipment use to the emissions response, including the electricity required to operate the equipment; $\psi_i^{p,f}$ is the response of fuel f relative to total electricity use; and Ξ_i^p is the response of total electricity use per dollar of revenue. The weight $\mathcal{W}_i^{p,f}$ converts relative-fuel responses into their contributions to pollutant p emissions, including the induced adjustment of process emissions through the outer production nest. The final term, Υ_i^p, is the revenue-scale response.

The propositions characterize responses to an infinitesimal increase in a positive permit price around a regular equilibrium. The focal regulated plant's equipment choice is interior. The sign results hold under the conditions stated in Appendix E.2.

Proposition 1 (Within-plant adjustment)**.** *A higher permit price τ_i^p has the following local*

effects on the three within-plant margins:

$$\Gamma_i^p < 0, \qquad \psi_i^{p,f} < 0, \qquad \Xi_i^p > 0.$$

That is, a higher pollution price increases the use of emission-control equipment, reduces fossil-fuel use relative to total electricity use, and raises total electricity use per dollar of revenue.

The first proposition maps to the observed changes in the use of emission-control equipment, fuel composition, and electricity use intensity. Higher electricity use per dollar of revenue can shift emissions to power generation.

Proposition 2 (Market-mediated leakage)**.** *For a regulated plant, a higher pollution price τ_i^p raises its unit cost and price and lowers its revenue. At an unregulated plant in district l serving a common destination market, revenue and pollutant p emissions rise when substitution toward its variety outweighs any contraction in demand for its firm composite:*

$$\Upsilon_i^p < 0, \qquad \frac{d \ln R_l}{d \ln \tau_i^p} > 0, \qquad \frac{d \ln s_l^p}{d \ln \tau_i^p} > 0.$$

The second proposition explains how unregulated plants can scale up through leakage and make a negative emissions difference-in-differences coefficient larger in magnitude. A higher pollution price raises the regulated plant's cost and price, shifting demand toward unregulated plants serving the same markets. For an unregulated affiliate, two effects work in opposite directions: it gains sales from the firm's regulated plants, but the firm as a whole loses sales to competitors. Since substitution across plants within a firm is stronger than substitution across firms ($\eta > \sigma$), the first effect dominates, and the affiliate expands.

Unregulated plants in competing firms also gain sales. These plants expand production without changing emissions per unit of output, so their emissions rise. If regulated emissions fall and the unregulated group includes plants receiving this additional demand, the emissions difference-in-differences is more negative than the regulated group's response alone: it combines a reduction at regulated plants with an increase at unregulated plants.

These propositions will guide the quantitative analysis: Proposition 1 links the observed abatement responses to the model parameters estimated in Section 5, while Proposition 2 motivates the decomposition of relative emissions responses in Section 6.

4.3 Power Plants

Power plants are classified by fuel type $g \in$ {Clean, Coal, Gas, Oil}, where Clean aggregates the non-fossil and new-and-renewable sources, such as hydro, nuclear, solar, and wind.[21] Unlike manufacturing plants, generators cannot switch fuels in the short run and operate under centralized dispatch rather than monopolistic competition.

There is a fixed mass of power plants of each fuel type g in a district i, $M_{i,g}$. A power plant in this group draws productivity (inverse fuel intensity) z. These plants generate electricity up to their capacity. To reduce dimensionality, I assume an engineering relationship between capacity k and productivity z, reflecting that larger and more efficient plants tend to have higher capacity (see Christensen and Greene 1976).

4.3.1 Electricity Generation and Emissions

Let y denote net electricity output assigned to a power plant in district i, with fuel type g and productivity z. Conditional on dispatched output y, the plant minimizes its total variable generation cost by choosing fuel use and the use of emission-control equipment:

$$V_{i,g}(y,z) = \min_{e_g; a \in [0,1)} \left\{ p^g e_g + \sum_{p \in \Lambda_{i,g}(z)} \tau_i^p m_{i,g}(z)\, s_{i,g}^p(y,z) \right\}$$
$$\text{s.t. } y = (1-a) \min\left\{ k_{i,g}(z), z e_g \right\},$$
$$s_{i,g}^p(y,z) = (1-a)^{\frac{1}{\xi^p}} \kappa_g^p e_g,$$

where $0 < \xi^p < 1$. p^g denotes the fuel g price for power plants, which is zero for non-fossil-fuel power plants. κ_g^p denotes emission factors of pollutant p in the power generation sector with energy source g. As in manufacturing, the regulation is applied only to plants whose unregulated emissions meet or exceed the thresholds: $\hat{s}_{i,g}^p(z) \geq \bar{T}^p$. Because pollution prices τ_i^p vary by pollutant and district, the set of regulated pollutants $\Lambda_{i,g}(z)$ depends on district i, fuel type g, and productivity z. To align with the policy, I allow non-fossil-fuel generators using polluting energy sources, such as biofuel and waste, to emit regulated pollutants. As in manufacturing, permits are allocated in lump sum, based on the pre-regulation emissions discounted by the cap changes. Let $A_{i,g}^p(z)$ denote the initial allocation of permits for pollutant p to the power plant. $m_{i,g}(z)$ captures plant-specific differences in effective compliance costs and is common across pollutants.

At a given output level, power plants can abate only through emission-control equip-

[21] Before the regulation, Clean was mostly nuclear in South Korea.

ment. $a_{i,g}(y,z)$ is the share of gross generation used to operate emission-control equipment. For positive dispatched output, feasibility implies $0 \leq a_{i,g}(y,z) < 1$. The parameter ξ^p governs the curvature of the relationship between equipment use and power-sector emissions and hence the effectiveness of abatement in accommodating additional electricity demand. y is net electricity supplied to the grid: the plant generates electricity using a CRS technology and uses a portion of it for abatement. Unlike in manufacturing, finite capacity $k_{i,g}(z)$ breaks the equality between marginal and unit cost. While capacity is slack, power generation is CRS and marginal generation cost equals unit cost. Once capacity binds, marginal generation cost is output-dependent and differs from average unit cost.

The cost minimization problem can be solved as follows. For a given equipment choice a, cost-minimizing fuel use is $e_g = y/(z(1-a))$. Substituting this expression and the emissions equation into the objective reduces the problem to choosing equipment intensity subject to the capacity constraint $y \leq (1-a)\,k_{i,g}(z)$. Let $a^0_{i,g}(z)$ denote the cost-minimizing emission-control equipment choice when the capacity constraint is slack. If $p^g = 0$ and $\tau_i^p \kappa_g^p > 0$ for at least one $p \in \Lambda_{i,g}(z)$, the slack-capacity cost expression, considered without the capacity constraint, has no finite minimizer: with free fuel, permit costs approach zero as $1-a$ approaches zero. I therefore define the limiting slack-capacity choice by $1 - a^0_{i,g}(z) = 0$. If the slack-capacity choice cannot accommodate dispatched output, capacity binds and $1 - a = y/k_{i,g}(z)$. The cost-minimizing equipment choice is therefore:

$$1 - a_{i,g}(y,z) = \max\left\{1 - a^0_{i,g}(z), \frac{y}{k_{i,g}(z)}\right\}. \tag{8}$$

When $1 - a^0_{i,g}(z) = 0$, every positive dispatched output lies in the capacity-bound branch. If $p^g = 0$ but $\tau_i^p \kappa_g^p = 0$ for every p, I set $a^0_{i,g}(z) = 0$. An unregulated plant likewise has $a^0_{i,g}(z) = 0$.

Where capacity is slack, substituting the optimal equipment choice into the objective gives the effective cost per unit of fuel:

$$\widetilde{p}_{i,g}(y,z) = p^g + \sum_{p \in \Lambda_{i,g}(z)} \tau_i^p m_{i,g}(z) \left(1 - a^0_{i,g}(z)\right)^{\frac{1}{\xi^p}} \kappa_g^p.$$

Differentiating minimized total cost $V_{i,g}(y,z)$ with respect to dispatched output y gives

the marginal generation cost:

$$\nu_{i,g}(y,z) = \underbrace{\frac{1}{\left(1-a^0_{i,g}(z)\right)}}_{\text{Emission-control cost}} \left(\underbrace{\frac{\widetilde{p}_{i,g}(y,z)}{z}}_{\text{Factor cost}} \right),$$

if $0 < y \leq \left(1-a^0_{i,g}(z)\right) k_{i,g}(z)$ and

$$\nu_{i,g}(y,z) = \underbrace{\frac{1}{z} \sum_{p \in \Lambda_{i,g}(z)} \frac{\tau_i^p m_{i,g}(z) \kappa_g^p}{\xi^p} \left(\frac{y}{k_{i,g}(z)} \right)^{\frac{1}{\xi^p}-1}}_{\text{Marginal pollution cost}},$$

if $\left(1-a^0_{i,g}(z)\right) k_{i,g}(z) < y \leq k_{i,g}(z)$. At zero output, I set $V_{i,g}(0,z) = e_{i,g}(0,z) = s^p_{i,g}(0,z) = 0$ and define $\nu_{i,g}(0,z)$ by its right-hand limit.

In the capacity-bound range, gross generation is fixed at $k_{i,g}(z)$, so fuel use is fixed at $e_{i,g} = k_{i,g}(z)/z$. Additional net output is obtained by reducing emission-control use rather than by using more fuel. The fuel-cost component is therefore fixed with respect to y, so p^g does not enter the marginal generation cost. Appendix E provides the first-order and boundary conditions.

4.3.2 Electricity Market

Cost changes reallocate production across plants in both sectors, but market structure affects how these changes translate into prices. For tractability, I model the government-controlled wholesale and retail markets as a single market without the government's intermediary role. I retain two key features. First, generators are dispatched in merit order (merit-order dispatch, MOD): incremental generation is dispatched in ascending order of the plant-level marginal-cost schedules, up to each plant's physical capacity. Second, under this rule, the system marginal price (SMP) equals the marginal generation cost of the final generation increment needed to meet aggregate demand, yielding a uniform nationwide electricity price $p^e = \nu^*$, where ν^* is the cost of the marginal increment in the system-wide ordering of $\nu_{i,g}(y,z)$. See Appendix A.2 for more details.[22]

[22]Transmission and distribution constraints are abstracted from the analysis, given Korea's near-100% grid coverage. For such features, see Arkolakis and Walsh (2023). Also, electricity prices in Korea vary only over time (e.g., through time-of-use pricing) and across broad user categories, such as residential or industrial consumers.

Electricity demand consists of manufacturing demand, ED, including equipment operation, and residual demand outside manufacturing, ED_R. The electricity market-clearing condition is:

$$\sum_g \sum_i M_{i,g} \mathbb{E}_z \left[y_{i,g}(z) \right] = ED + ED_R. \tag{9}$$

The capacity constraint is essential: without it, zero-cost non-fossil-fuel plants could supply arbitrarily large quantities and drive the SMP to zero. With finite clean generation capacity, these plants are dispatched first but can meet demand only up to their installed capacity, so additional demand is more likely to be met by fossil-fuel plants, creating unintended emissions. Where dirty power plants are located matters in evaluating emissions shifts into the power sector. However, unlike in manufacturing, within-firm and geographic pollution shifts are constrained by the dispatch order, because centralized dispatch determines reallocation across generators.

Equipment use changes generators' marginal-cost schedules, but p^e is determined by the marginal cost of the last generation increment needed to meet total electricity demand under merit-order dispatch. Regulation therefore affects the electricity price only through the cost of this marginal increment, so regulated manufacturers' permit-inclusive fossil-fuel prices can rise relative to the electricity price. Industrial electrification can thus arise from differences in abatement technology and productivity across sectors.

After dispatch, the profit of a power plant can be expressed as:

$$\pi_{i,g}(z) = p^e y_{i,g}(z) - \int_0^{y_{i,g}(z)} v_{i,g}(u,z)\, du + \sum_{p \in \Lambda_{i,g}(z)} \tau_i^p A_{i,g}^p(z).$$

4.4 Households

Representative households consume final goods, Q_n, and experience pollution damage, SCE_n. For consumption, there are three layers. Plant varieties are combined into firm goods, firm goods are combined into sectoral goods, both using Dixit–Stiglitz preferences (Dixit and Stiglitz 1977), and sectoral goods are then combined across sectors under Cobb–Douglas preferences. In sector j, let σ_j denote the elasticity of substitution across firm goods ω and η_j denote the elasticity of substitution across plant varieties within firms $\tilde{\omega}$. The latter CES aggregate captures the stronger within-firm leakage than across-firm leakage within industries. Aggregators and corresponding ideal price indices are given by:

$$Q_n = \prod_j Q_{n,j}^{\gamma_j^C}, \qquad P_n = \prod_j \left(\frac{P_{n,j}}{\gamma_j^C} \right)^{\gamma_j^C}, \tag{10}$$

$$Q_{n,j} = \left(\sum_i \int_{\omega \in \Omega_{in,j}} Q_{in,j}^{\frac{\sigma_j - 1}{\sigma_j}} (\omega)\, d\omega \right)^{\frac{\sigma_j}{\sigma_j - 1}}, \; P_{n,j} = \left(\sum_i \int_{\omega \in \Omega_{in,j}} P_{in,j}^{1-\sigma_j} (\omega)\, d\omega \right)^{\frac{1}{1-\sigma_j}},$$

$$Q_{in,j}(\omega) = \left(\sum_l \int_{\tilde{\omega} \in \Omega_{iln,j}(\omega)} q_{iln,j}^{\frac{\eta_j - 1}{\eta_j}} (\tilde{\omega})\, d\tilde{\omega} \right)^{\frac{\eta_j}{\eta_j - 1}}, \; P_{in,j}(\omega) = \left(\sum_l \int_{\tilde{\omega} \in \Omega_{iln,j}(\omega)} p_{iln,j}^{1-\eta_j} (\tilde{\omega})\, d\tilde{\omega} \right)^{\frac{1}{1-\eta_j}},$$

where $\eta_j > \sigma_j > 1$ and $\sum_j \gamma_j^C = 1$. $\Omega_{in,j}$ is the set of goods sold in market n by firms headquartered in district i, while $\Omega_{iln,j}(\omega)$ is the set of plant varieties of firm ω headquartered in district i, produced in district l, and sold in market n in sector j. γ_j^C denotes the sectoral expenditure share. The corresponding demands at the firm and plant levels are:

$$\begin{aligned} Q_{in,j}(\omega) &= \left(\frac{P_{in,j}(\omega)}{P_{n,j}} \right)^{-\sigma_j} Q_{n,j}, \\ q_{iln,j}(\tilde{\omega}) &= \left(\frac{p_{iln,j}(\tilde{\omega})}{P_{in,j}(\omega)} \right)^{-\eta_j} Q_{in,j}(\omega) \\ &= \left(\frac{p_{iln,j}(\tilde{\omega})}{P_{in,j}(\omega)} \right)^{-\eta_j} \left(\frac{P_{in,j}(\omega)}{P_{n,j}} \right)^{-\sigma_j} Q_{n,j}. \end{aligned} \tag{11}$$

I construct a pollution damage index for district n, SCE_n, as:

$$SCE_n = \frac{L_n}{\sum_m L_m} \times \sum_{p \in \{\mathrm{NO}_x, \mathrm{SO}_x, \mathrm{TSP}\}} SC^p \times S_n^p, \tag{12}$$

where S_n^p denotes the total emissions of pollutant p in district n. In the model, aggregate emissions are defined as $S_n^p = \sum_j S_{n,j}^p + \sum_g S_{n,g}^p$, where $S_{n,j}^p$ and $S_{n,g}^p$ denote total emissions of p in district n from manufacturing sector j and from power generation using fuel g, respectively. SC^p denotes the social cost per unit of pollutant p and serves as an aggregation weight reflecting differences in harm across pollutants. I report percentage changes in real consumption and this damage index separately, rather than aggregating them into a single welfare measure or calculating monetary net benefits. The damages from local air pollution also vary across districts because exposure is greater in more populated districts. To capture spatial variation in exposure, I use district labor endowment, L_n, as a proxy for population and weight pollution damages by $L_n / (\sum_m L_m)$. I assume pollution is contained within each district and abstract from atmospheric dispersion.

The representative household in district n faces the budget constraint:

$$P_n Q_n = w_n L_n + \Pi_n + r K_n, \tag{13}$$

where K_n and L_n are the capital and labor endowments. The first term is labor income. Labor is mobile across sectors but immobile across districts, being supplied at wage w_n. Households cannot respond to local pollution or income changes by moving across districts.[23] The second term aggregates profits of manufacturing firms and power plants headquartered in district n. Permit rents do not affect marginal plant decisions. The value of freely allocated permits is included in firm profits and accrues to households in firms' headquarters districts. The third term is household capital income rK_n. Capital services are fully mobile across sectors and districts and supplied perfectly elastically at a common rental rate r, taken as given and normalized to one.

4.5 Equilibrium

This section derives expressions for aggregate variables and defines the equilibrium of the model. Manufacturing firm productivity draws ϕ are independently and identically distributed (i.i.d.) according to a Pareto distribution with district–sector shifter $T_{i,j}$ and shape parameter $\zeta_j > 0$. Plant productivity draws z are i.i.d. and $\ln z$ is normally distributed with mean $-\left(\sigma_j^z\right)^2/2$ and standard deviation σ_j^z. Power plant productivity draws z follow a Pareto distribution with district–fuel shifter $T_{i,g}$ and shape parameter $\zeta_g > 0$.

Market clearing conditions hold in electricity, labor, product, and permit markets. The electricity market clearing is explained in Equation (9). I abstract from fuel extraction, reflecting Korea's status as a small open economy importing approximately 95% of its coal, gas, and oil annually, with full imports of oil and gas over the past decade. Hence, supplies of fossil fuels are assumed perfectly elastic. Fuel markets are perfectly competitive, and fuels are freely traded across districts.

The labor market in each district clears:

$$L_l = \sum_{i,j} M_{i,j} \mathbb{E}_{\phi,z} \left[\ell_{il,j} \left(\phi, z \right) \right], \tag{14}$$

where $\ell_{il,j} \left(\phi, z \right)$ denotes the labor demand at the plant level.

Households allocate income across sectors according to their Cobb–Douglas expendi-

[23] Because the policy targets short-lived local pollutants, I hold district populations fixed to isolate the short-run production and abatement responses to this place-based policy. Migration would introduce an additional adjustment margin, as households could relocate in response to changes in real income and pollution exposure.

ture shares. The product market clears in each district l and sector j:

$$X_{l,j} = \gamma_j^C \left(w_l L_l + r K_l + \Pi_l \right). \tag{15}$$

Lastly, each permit market clears:

$$0 \leq \tau_a^p \;\perp\; C_a^p - \sum_{l \in a} \left(\sum_j \bar{S}_{l,j}^p + \sum_g \bar{S}_{l,g}^p \right) \geq 0, \tag{16}$$

where policy area a is a set of regulated districts, and $\bar{S}_{l,j}^p$ and $\bar{S}_{l,g}^p$ denote the total amount of emissions of pollutant p in district l for regulated manufacturing plants in sector j and regulated power plants, respectively. C_a^p is the emissions cap for pollutant p in policy area a. For each pollutant, a common permit price clears the market within each policy area, so $\tau_l^p = \tau_a^p$ for all $l \in a$.

Let $\mathcal{P}^M$ and $\mathcal{P}^E$ denote the collections of manufacturing and power-generation policy functions, respectively:

$$\mathcal{P}^M \equiv \left\{ k_{il,j}, e_{il,j}, \ell_{il,j}, \left\{ e_{il,j}^f \right\}_{f \in \mathbb{F}}, c_{il,j} \right\}_{i,l,j},$$
$$\mathcal{P}^E \equiv \left\{ y_{l,g}, e_{l,g}, a_{l,g} \right\}_{l,g}.$$

Function arguments are suppressed for brevity. Then, the spatial equilibrium of the economy across sectors, fuel types, pollutants, and districts is defined as follows.

Equilibrium Definition

Given capital and labor endowments $\{K_l, L_l\}$, manufacturing headquarters masses $\{M_{i,j}\}$, power plant masses $\{M_{l,g}\}$, a common rental rate r, non-electric manufacturing fuel prices $\{p^f\}_{f \in \mathbb{F} \setminus \{e\}}$, power plant fuel prices $\{p^g\}$, pollution caps $\{C_a^p\}$, and residual electricity demand ED_R, an equilibrium is defined with prices for sectoral goods $\{P_{l,j}\}$, wages $\{w_l\}$, an electricity price p^e, permit prices $\{\tau_l^p\}$, policy functions of manufacturing plants $\mathcal{P}^M$, policy functions of power plants $\mathcal{P}^E$, and market expenditures $\{X_{l,j}\}$ such that:

1. Policy functions solve the cost-minimization problems of manufacturing and power plants and the pricing problems of manufacturing firms, and power plant output follows merit-order dispatch.

2. The price index equations hold (Equation (10)).

3. Electricity, labor, product, and permit markets clear (Equations (9), (14), (15), and (16)).

Finally, the model features equilibrium channels linking abatement to production and emissions reallocation. In manufacturing, nested CES demand allows production to shift within and across manufacturing firms. Multiregional production frictions and iceberg trade costs shape reallocation across districts. In power generation, there is no leakage within firms because there is a centralized dispatch for a homogeneous good, electricity. Electricity demand links manufacturing abatement to emissions from power generation.

In the policy exercises, I hold wages and incomes at their unregulated values, without imposing labor market clearing conditions. Thus, consumption changes reflect changes in the price index. Plants adjust their input demands and production under the same optimality conditions. This industry equilibrium approach focuses on product, electricity, and pollution market responses (see Chen et al. 2025 and Curtis et al. 2026 for related work).

See Appendix G for the solution algorithm.

5 Calibration and Estimation

This section describes the calibration and estimation, presents the parameter estimates, and evaluates model fit. The model is implemented with $J = 10$ manufacturing sectors and $N = 250$ districts corresponding to Korea's administrative districts. All monetary values are expressed in thousands of U.S. dollars (USD), the energy unit is toe, and emissions are measured in metric tons.

I quantify the model in two steps. Pre-regulation data determine firm and plant productivity parameters, Cobb–Douglas factor shares, demand and fuel-substitution elasticities, and baseline emissions coefficients. Conditional on this baseline and observed permit prices, I then estimate equipment-effectiveness parameters, the physical electricity requirement of manufacturing equipment, and parameters governing source- and process-specific effective compliance costs, using cap-reduction benchmarks and the reduced-form responses.

5.1 Calibration

I use averages over 2016–2019 to reduce the influence of year-specific shocks. I calibrate each pollutant-specific regulatory threshold, $\bar{T}^p$, using a proxy for baseline regulatory

coverage: the ratio of the aggregate emissions cap in 2020 to the emissions of regulated sectors in 2018 within regulated districts. I apply this proxy to manufacturing and power generation because these sectors account for most covered emissions. I pool modeled manufacturing and power plants in regulated districts and rank them by baseline emissions of each pollutant. I choose a common cutoff so that plants above it account for a share of baseline emissions as close as feasible to the target, subject to retaining at least two units of regulated plant mass in every permit market. The latter condition makes each market well-defined. The corresponding coverage targets are 87%, 63%, and 18% for NO_x, SO_x, and TSP, respectively. Among modeled manufacturing and power plants in regulated districts, the thresholds cover 87% of baseline NO_x emissions, 63% of SO_x, and 21% of TSP.

Capital and labor endowments, K_l and L_l, are constructed from the MMS and SBR as annual averages of tangible capital and employment. For labor, I use the number of regular employees whose tasks are either in mining or manufacturing. The numbers of firms and plants across sectors, fuel types, and districts are taken from the MMS and SBR.[24] I generalize the plant-mass weights while retaining the model's production, demand, and pricing blocks. Let $M_{il,j}$ denote the observed mass of plants in district l and sector j owned by firms headquartered in district i. For $M_{i,j} > 0$, I use $M_{il,j}/M_{i,j}$ as the plant mass per firm in district l, held common across firm-productivity draws. This ratio weights each location's contribution to the firm price index and expected profit. Aggregate plant-level labor demand, electricity use, and emissions are weighted by $M_{il,j}$, whereas aggregation across firms retains $M_{i,j}$. Because disclosure restrictions prevent the release of firm-level plant networks, I use this common plant-mass profile as an approximation.

Fuel prices p^f and p^g are sourced from the Korea Energy Economics Institute. Iceberg trade costs $t_{ln,j}$ between production district l and destination n in sector j are parameterized as in Arkolakis and Walsh (2023). Sectoral final expenditure shares, γ_j^C, are based on the Input–Output Table compiled by the Bank of Korea in 2016. I estimate the across-firm substitution elasticities, σ_j, and calibrate the within-firm substitution elasticities, η_j, following Broda and Weinstein (2006).

For manufacturing, I jointly calibrate location productivity parameters $T_{i,j}$ and multiregional production frictions $\gamma_{il,j}$ using the plant-count-weighted, normalized distributions of annual average sales per active plant across headquarters sector–district cells and affili-

[24]Within each district, I allocate fossil power-plant counts across fuels using pre-regulation emissions shares. The allocation weights use each fuel's share of the combined mass of emissions. Because fuel-specific profits are unavailable, I assign the observed district-level mean profit of fossil power plants to each active fossil-fuel category. This procedure preserves district-level plant counts and total profits; the fuel-level allocation is an approximation rather than separately observed data.

ate headquarters–plant district dyads, respectively. I discipline σ_j^z using the sector-level mean of the largest plant sales share within a firm–year–sector, restricting both the data and the model to firms with at least two plants. Statistical disclosure restrictions permit the release of this sector-level moment and aggregate headquarters–plant margins, but not the underlying firm–plant incidence matrix. I approximate this moment using synthetic finite plant portfolios consistent with the released margins. These portfolios are used only to construct the concentration moment and do not determine the equilibrium allocation. See Appendix F.2 for more details.

For power generation, I calibrate $T_{i,g}$ using the firm-count-weighted, normalized distribution of annual average variable profits across fuel–district cells. To discipline relative scale across these blocks, I also target the ratio of total affiliate sales to total headquarters sales and the ratio of total power generation variable profits to total manufacturing sales. For power generation capacity, I assume a log-linear relationship between capacity and productivity, $\ln k_{i,g}(z) = \mathcal{A} \ln z$, which reflects that larger power plants are more productive. The parameter $\mathcal{A} = 0.12$ is calibrated using data on inputs, outputs, and capacity for fossil-fuel power plants. This positive capacity–productivity relationship is consistent with evidence of heterogeneous plant productivity and economies of scale in Korean power generation (Heshmati et al. 2014; Oh and Lee 2016). I use a constant residual electricity demand outside manufacturing to match the pre-regulation electricity-to-gas price ratio. This is because the MOD and SMP determine the electricity price based on total electricity demand.

I then calibrate physical emissions coefficients as ratios of observed pre-regulation emissions to the corresponding model-implied output or fuel use. Output-linked manufacturing emissions, including process emissions and emissions from fuels not separately represented in the model, are pooled by pollutant and sector; combustion emissions from modeled manufacturing fuels are pooled by pollutant and fuel; power generation emissions are pooled by pollutant and generation fuel. These coefficients reproduce baseline emissions totals within the pooled groups. I do not require an exact match within every district, so location-specific emissions coefficients do not mechanically absorb regional differences in plant composition. This pooling is consistent with the emission factors given by the Ministry of Environment and preserves key emission sources within each manufacturing sector and power generation energy source.

5.2 Estimation

I first estimate production and demand parameters using pre-regulation data, then use the simulated method of moments to estimate abatement and effective compliance costs.

5.2.1 Demand and Production Elasticities

I estimate the baseline production parameters using 2016–2019 data, before the cap-and-trade expansion. In the model, the equipment-electricity and emissions-cost terms associated with this expansion are inactive in the baseline and therefore omitted from first-stage estimation. I calibrate the Cobb–Douglas factor shares and estimate the CES fuel technology and firm-level demand elasticity. I allow input prices to vary across plants in estimation. As in the quantitative model, I use revenue net of material expenditures and assume that it satisfies the revenue equation implied by a CES restricted production technology over capital, labor, and energy. Estimation is conducted separately by sector; I suppress plant, district, and sector indices below.

First, I calibrate the outer factor shares from input expenditures. Let $\mathcal{S}^K$ be the sum of capital expenditure, $\mathcal{S}^L$ be the sum of labor expenditure, and $\mathcal{S}^E$ be the sum of fuel expenditure, and define $\mathcal{S} = \mathcal{S}^K + \mathcal{S}^L + \mathcal{S}^E$. Capital expenditure is the rental or user cost of capital services, not the value of the capital stock. The Cobb–Douglas first-order conditions imply:

$$\alpha^K = \frac{\mathcal{S}^K}{\mathcal{S}}, \qquad \alpha^E = \frac{\mathcal{S}^E}{\mathcal{S}}, \qquad \alpha^L = 1 - \alpha^K - \alpha^E,$$

where the sums include all plant-year observations in the sector.

Second, I control for the transmission of Hicks-neutral productivity ϕz into input choices, following Doraszelski and Jaumandreu (2013) and Doraszelski and Jaumandreu (2018). I use the electricity first-order condition to estimate the fuel-substitution elasticity λ, the relative fuel weights α^f/α^e, and the firm-level demand elasticity σ. Define $\widetilde{R} = \alpha^E R/\left(p^e e^e\right)$. The estimating equation is:

$$\ln \widetilde{R} = \ln\left(\frac{\sigma}{\sigma-1}\right) + \ln\left[1 + \sum_{f \neq e} \frac{\alpha^f}{\alpha^e}\left(\frac{e^f}{e^e}\right)^{\frac{\lambda-1}{\lambda}}\right] + u. \tag{17}$$

Here, u is a log-revenue measurement disturbance. This equation uses observed fuel quantities measured in tons of oil equivalent. I estimate it by nonlinear least squares (NLLS), following Grieco et al. (2016) and Murray Leclair (2024) under the maintained

conditional mean restriction $\mathbb{E}[u \mid \mathcal{I}] = 0$, where $\mathcal{I}$ contains the information used for input choices, including input prices and the realized productivity and demand conditions. Appendix F provides more details on the derivation. The normalization $\sum_{f \in \mathbb{F}} \alpha^f = 1$ gives $\alpha^e = \left[1 + \sum_{f \neq e} \left(\alpha^f / \alpha^e\right)\right]^{-1}$ and $\alpha^f = \alpha^e \left(\alpha^f / \alpha^e\right)$ for $f \neq e$.

Conditional on the calibrated outer energy share α^E, identification of λ and the relative fuel weights relies on sufficiently rich variation in fuel quantities relative to electricity use. The level of revenue relative to electricity expenditure identifies the markup, and hence σ, under the stated conditional mean restriction. In general, the demand elasticity σ is not cleanly pinned down by the production side alone in a revenue specification (De Loecker 2011). Under Dixit–Stiglitz demand, however, σ can be estimated jointly with the production parameters without observing output prices or quantities.

The approach requires plants to take input prices as given and to adjust energy inputs flexibly within the period so that the electricity first-order condition holds. Unobserved input-price differences across plants need not invalidate this approach (Grieco et al. 2016). Such differences may generate the variation in fuel composition used for identification. The concern is whether procurement arrangements or other input-market frictions invalidate these optimality conditions or the conditional mean restriction. Quantity-dependent prices, adjustment costs, binding fuel-access constraints, and measurement error in fuel quantities or expenditures can do so; differences in price levels alone do not. These concerns are particularly relevant for large polluting plants, whose procurement arrangements may differ systematically from those of other plants. To assess sensitivity to firm organization and associated procurement arrangements, I re-estimate the same specification using only single-plant firms. The estimated fuel-substitution and firm-level demand elasticities remain similar; see Figure A.10.

5.2.2 Productivity Dispersion Parameters

I focus on the pre-regulation period using the 2016 MMS and SBR for manufacturing and the 2015 Economic Census and SBR for power generation, respectively. In the manufacturing sector, I exploit the properties of the Pareto and log-normal distributions and the CES demand to estimate the shape parameter ζ_j, focusing on the right tail of the sales distribution. If firm productivity follows a Pareto distribution with shape parameter ζ_j and log plant productivity follows a normal distribution with mean $-\left(\sigma_j^z\right)^2/2$ and standard deviation σ_j^z, then the tail index is $\zeta_j / \left(\sigma_j - 1\right)$.[25] Similarly, in the power generation sector,

[25] In the pre-regulation model, holding the plant portfolio and district-sector characteristics fixed, $p_{iln,j}(\phi, z) \propto (\phi z)^{-1}$ and $P_{in,j}(\phi) \propto \phi^{-1}$. The nested CES revenue equation therefore implies

I estimate the shape parameter ζ_g using the relationship between productivity and capacity. Given the log-linear capacity function, if productivity follows a Pareto distribution with shape parameter ζ_g, then the tail index is $\zeta_g / \mathcal{A}$, since capacity is likely to bind for plants with high productivity.

5.2.3 Simulated Method of Moments

Pre-policy equipment operating hours provide a proxy for differences in compliance conditions across emissions sources and production processes. Let h index a plant in manufacturing sector or generation fuel group j. Define $B_h^U = \sum_p s_h^{p,U}$ as its total emissions in the unregulated baseline, and let ω_h denote its plant-mass and numerical weight. Define the group mean as $\bar{B}_j^U = \sum_{h \in j} \omega_h B_h^U / \sum_{h \in j} \omega_h$. I construct an index of compliance conditions according to:

$$H_h = \min \left\{ 8760, A_j \left(1 + \frac{B_h^U}{\bar{B}_j^U} \right)^{\rho_j^H} \right\},$$

where $\rho_j^H \geq 0$. The upper bound of 8,760 is the total number of hours in a non-leap year. A_j ensures that the weighted group mean equals the observed pre-policy mean annual equipment operating hours:

$$\frac{\sum_{h \in j} \omega_h H_h}{\sum_{h \in j} \omega_h} = \bar{H}_j^{\text{Data}}.$$

The manufacturing means are sector-specific. For power generation, I use the pooled mean for all four fuel groups. Annual hours are daily operating hours multiplied by annual operating days; the group means use positive observations before 2020. When $\rho_j^H = 0$, the index is constant within the group. Motivated by the greater pre-policy equipment use among regulated plants in Table A.5, a positive ρ_j^H assigns higher index values to larger baseline emitters within each group, subject to the hours ceiling and the observed group mean.

The effective compliance cost multiplier is:

$$m_h = \lambda_H \left(1 + \frac{H_h}{H_{\text{Ref}}} \right)^{-\beta_H}.$$

The reference level is $H_{\text{Ref}} = 2{,}400$ hours. λ_H is a common cost scale and β_H governs the

$R_{il,j}(\phi, z) \propto \phi^{\sigma_j - 1} z^{\eta_j - 1}$. Since ϕ and z are independent and log-normal z has finite moments of every positive order, Breiman's theorem implies a sales-tail index of $\zeta_j / (\sigma_j - 1)$.

sensitivity of effective compliance costs to the index. Within manufacturing, the same plant-specific multiplier applies to process and combustion emissions. This preserves the sector's underlying emissions composition: the regulatory burden operates primarily through production where output-linked emissions dominate and through fossil-fuel use where combustion emissions dominate.

Then, I jointly estimate 22 parameters governing abatement technology and effective compliance costs. The seven technology parameters comprise three manufacturing equipment-effectiveness parameters α^p, three power-generation equipment-effectiveness parameters ξ^p, and the manufacturing equipment electricity requirement δ. In estimation, I restrict the manufacturing parameters to $\alpha^p \leq \ln 100 / \ln 2$, which limits direct emissions removal to 99% at full equipment intensity, $c = 1$. The remaining 15 parameters govern compliance costs: a common scale λ_H, a compliance-cost slope β_H, and within-group index slopes ρ_j^H for nine manufacturing sectors and four generation fuel groups. The mining slope is fixed at zero because mining has no covered baseline emissions in the model. I use 22 targets: 12 market-specific cap-reduction benchmarks and 10 reduced-form estimates. The response moments comprise six pollutant-specific emissions responses, equipment-use responses in manufacturing and power generation, the sales response of unregulated manufacturing affiliates, and the response of electricity's share in manufacturing energy use.

Having calibrated the model's emissions coefficients to observed pre-regulation emissions, I use observed permit prices to price these emissions in estimation rather than determining prices through permit-market clearing. Manufacturing and power generation account for most regulated plants and a large share of emissions from regulated sectors in regulated districts. I therefore treat the emissions levels implied by market-wide cap-reduction rates as approximate ceilings for the modeled sectors. The estimation criterion places a smaller penalty on reductions above these benchmarks than on equally sized shortfalls.

The model propositions guide the choice of moments for the joint estimation. Pollutant-specific emissions responses and equipment use jointly inform equipment effectiveness in both sectors, while the manufacturing electricity share helps discipline the equipment electricity requirement. Affiliate sales and the cap-reduction benchmarks inform the common compliance-cost parameters and within-group index slopes. Appendix F provides further details.

Table 4: Demand and production parameters

Industry	Demand elasticity σ	η	Fuel elasticity λ	Pareto shape parameter ζ
Manufacturing				
Mining/quarrying	1.37 (0.05)	9.35	1.80 (0.45)	0.52 (0.03)
Food/beverages/tobacco	1.28 (0.01)	15.92	1.93 (0.10)	0.43 (0.01)
Textiles/apparel/leather	1.24 (0.01)	12.35	2.59 (0.26)	0.32 (0.004)
Wood/paper/printing/media	2.65 (0.10)	11.27	2.46 (0.12)	2.43 (0.05)
Petro/chem/pharma/rubber/plastics	1.14 (0.003)	9.66	2.05 (0.08)	0.17 (0.001)
Non-metallic minerals	1.31 (0.01)	7.36	2.02 (0.11)	0.43 (0.01)
Basic/fabricated metals	1.46 (0.01)	11.21	2.28 (0.11)	0.57 (0.003)
Electronics/instruments/electrical	1.08 (0.001)	9.94	1.93 (0.18)	0.09 (0.001)
Machinery/transport (incl. auto)	1.25 (0.004)	17.63	1.83 (0.18)	0.29 (0.002)
Furniture/misc./repair	3.33 (0.24)	5.20	2.48 (0.18)	3.81 (0.04)
Power generation				
Fossil-fuel				0.23 (0.02)
Non-fossil-fuel				0.08 (0.01)

Notes: The table reports estimated and calibrated parameters. The first parameter column reports estimates of σ_j; the second reports externally calibrated values of η_j; the third reports estimates of λ_j; and the fourth reports estimates of ζ_j and ζ_g. Standard errors are in parentheses.
Sources: Statistics Korea; Broda and Weinstein (2006); World Bank World Integrated Trade Solution concordances.

5.3 Results

Table 4 reports the first-stage parameter estimates; Table A.27 reports the corresponding factor shares and CES fuel weights. The estimated firm-level demand elasticities and calibrated within-firm elasticities satisfy $1 < \sigma_j < \eta_j$, so substitution across plants within a firm exceeds substitution across firms. The Cobb–Douglas outer nest fixes the elasticity of substitution across capital, energy, and labor at one. The inner CES nest allows greater substitution across fuels through the estimated parameter $\lambda_j > 1$, consistent with the maintained nesting restriction in the production function. Estimated fuel substitution is greater in light manufacturing, such as textiles, apparel, and leather. This pattern is consistent with the more specific fuel requirements of heavy industries, such as machinery and transport equipment.

The estimated fuel-substitution elasticities are broadly consistent with the values used in related studies. For comparison, Murray Leclair (2024) estimates an elasticity of 2.2 for the Indian steel industry. My estimates allow this substitution margin to vary across Korean manufacturing industries. The productivity shape parameters ζ_j satisfy $\zeta_j > \sigma_j - 1$

Table 5: Equipment-effectiveness parameters

	α^{NO_x}	α^{SO_x}	α^{TSP}
Manufacturing	6.51	6.34	4.37

	ξ^{NO_x}	ξ^{SO_x}	ξ^{TSP}
Power generation	0.31	0.10	0.87

Notes: The table reports equipment-effectiveness parameters estimated by the simulated method of moments.
Sources: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea.

for finite profits in manufacturing. Non-fossil-fuel power generation exhibits greater dispersion than fossil-fuel generation, reflecting input specificity and heterogeneity across its sources (Arkolakis and Walsh 2023).

Table 5 reports the equipment-effectiveness estimates. The manufacturing parameters α^p govern pollutant removal at a given equipment intensity, separately from factor and fuel substitution; the power generation parameters ξ^p govern the corresponding generation sector equipment technology. Under the actual policy, emission-control equipment accounts for 0.1% of manufacturing electricity use, implying a small additional electricity burden embedded in δ.

The common compliance-cost slope is $\beta_H = 36.9$, so modest differences in the index can generate substantial differences in effective compliance costs. Observed group-average hours generate differences across industries, while ρ_j^H governs how the index varies with baseline emissions within each industry. This gradient is largest in non-metallic minerals: $\rho_6^H = 0.004$, compared with 0.001 in six other manufacturing industries. Within non-metallic minerals, the plant-specific multiplier falls as low as 65.0% of the multiplier evaluated at the industry's mean operating hours. Thus, the estimated specification allows both industry-level differences and lower effective burdens for larger baseline emitters within an industry.

5.3.1 Model Fit

Targeted Moments The first-stage calibration closely reproduces the distributions of manufacturing sales and power generation variable profits (Table 6 and Figure A.11). The ratio of total affiliate sales to headquarters sales is 0.25 in the model and 0.26 in the data. The ratio of total power generation variable profits to total manufacturing sales is 0.02 in the model and 0.02 in the data. The pre-regulation electricity price is three times the gas price, measured per unit of energy.

Table 6: Model fit: production and emissions distributions

Moment	Used in estimation	Spearman correlation	Pearson correlation	N
Pre-regulation moments				
Manufacturing HQ sales	Yes	0.85	0.85	2,105
Manufacturing affiliate sales	Yes	0.82	0.84	2,986
Manufacturing maximum sales share	Yes	1.00	1.00	10
Power generation variable profits	Yes	0.94	0.94	278
Emissions distribution	Totals only	0.78	0.78	747
Post-regulation moments				
Emissions distribution	No	0.77	0.77	750

Notes: The table summarizes model fit for spatial distributions. HQ denotes headquarters. Spearman correlations measure agreement in the rank ordering of the data and model shares, whereas Pearson correlations measure the linear association between their logs. For the maximum sales share, Spearman and Pearson correlations are computed for ranks and levels, respectively, across the sector-level means and are weighted by sector-level HQ counts. The units of observation are sector–HQ district cells for manufacturing HQ sales; sector–HQ district-affiliate district cells for affiliate sales; fuel–district cells for power generation variable profits; and pollutant–district cells for emissions distributions. Baseline emissions totals are matched within pooled pollutant–source groups, but the district-level distribution is not directly targeted.
Sources: Ministry of Environment, Statistics Korea.

These first-stage moments discipline the baseline economic structure underlying the reallocation analysis. Manufacturing sales across headquarters and plant locations constrain productivity and multiregional production frictions, which shape relative delivered costs and the allocation of sales across plants and districts. Power generation profit patterns and the electricity-price target discipline the baseline electricity sector. The strong fit to the spatial distribution of affiliate sales is particularly relevant: it anchors the initial allocation of manufacturing activity from which the model predicts reallocation following regulation.

In the second stage, the model closely reproduces the targeted emissions responses. The model reproduces the signs of all 10 response targets, with nine lying within their empirical 95% confidence intervals. All six pollutant-specific coefficients lie within 0.7 empirical standard errors of their data counterparts. Averaging across pollutants, the manufacturing coefficient is −0.41 in the model and −0.43 in the data; the corresponding power generation coefficients are −0.23 and −0.22. The model also generates increased equipment intensity in both sectors, an increase in electricity's share of manufacturing energy use, and sales growth at unregulated affiliates. At observed permit prices, emissions reductions fall short of the cap-reduction benchmarks in only two of the 12 permit markets: SO_x in the Seoul Capital Area and the Central Area.

Some response magnitudes remain imperfectly matched. For example, manufacturing equipment intensity rises by 0.33, compared with a target of 0.17, while power generation intensity rises by 0.02, compared with a target of 0.90 under the maintained proxy mapping. The model measures continuous equipment intensity, whereas the empirical outcome is update-weighted log operating hours. This measurement difference may contribute to the discrepancy.

Untargeted Moments The model reproduces the broad spatial pattern of emissions before and after regulation. Across district-pollutant cells, both the Spearman and log Pearson correlations with the data are 0.78 before regulation and 0.77 afterward (Table 6 and Figure A.12). This set of moments is important because it relates to the geographic distribution of emissions, which is the key object of the model. The model reproduces the direction of the initial sales decline and negative employment responses in the event study.

6 Counterfactual Policy Experiments

I use the model to interpret the reduced-form estimates and compare consumption and pollution damage under the actual policy and nationwide alternatives. The unregulated baseline removes the modeled cap-and-trade incentives, setting $\tau_l^p = 0$ for all pollutants and districts. The actual-policy benchmark is the equilibrium evaluated at observed market-specific permit prices.

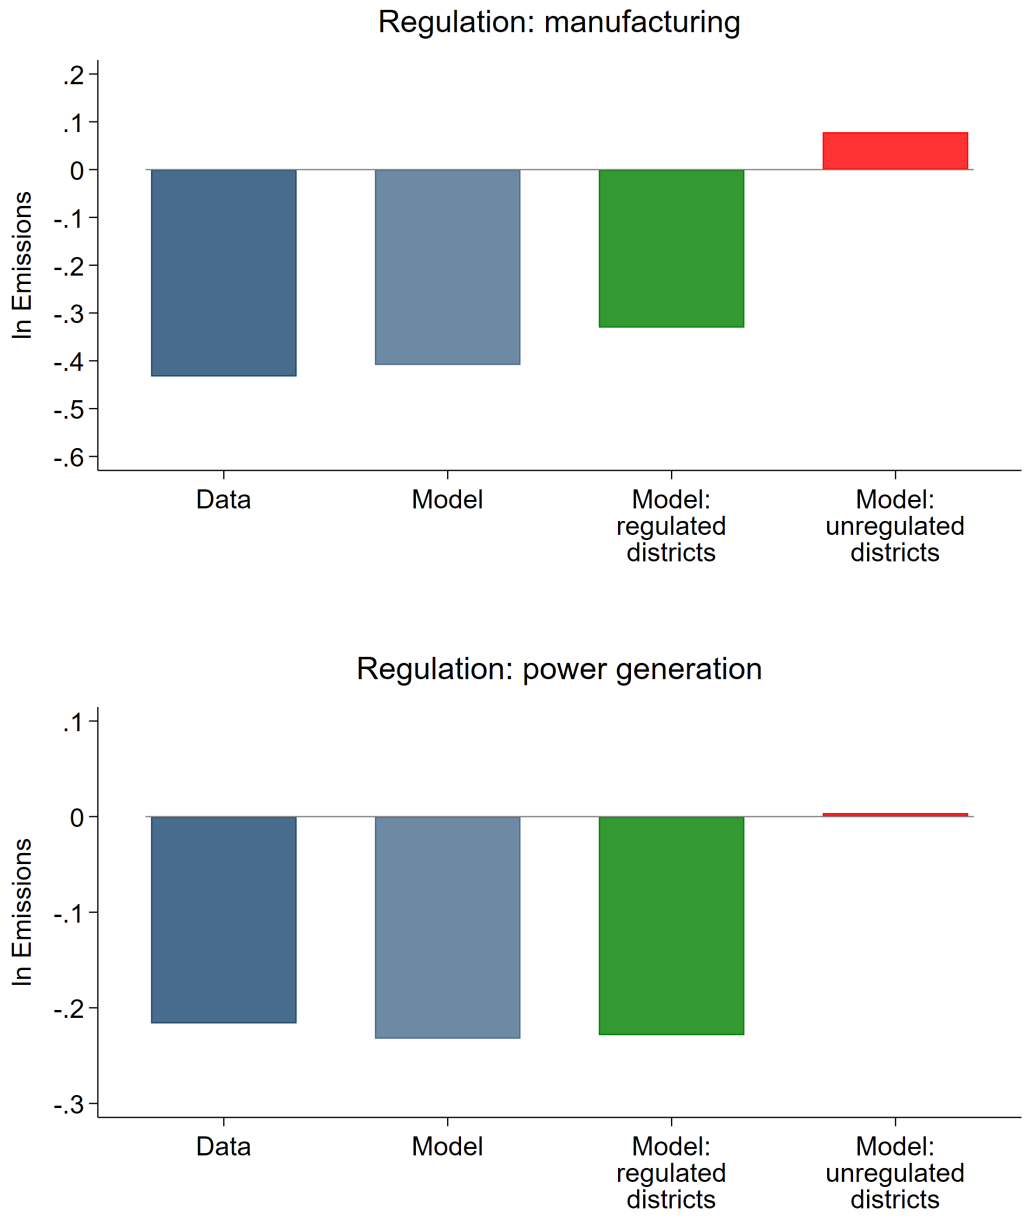
6.1 Total Volume Control of Pollutants at Place of Business

This section provides quantitative results related to the actual policy.

6.1.1 Reduced-Form Interpretation

I decompose the model-implied difference-in-differences coefficient into contributions from regulated and unregulated districts. In each panel of Figure 6, the first two columns show the data and model estimates for the corresponding sector: the model estimate closely matches the data estimate. The last two columns decompose the simulated difference-in-differences coefficient into regulated- and unregulated-cell contributions. Emissions reductions in regulated districts account for 80.9% of the model-implied manufacturing emissions difference-in-differences estimate, while emissions increases in unregulated

Figure 6: Decomposition of emissions responses



Notes: The averages of difference-in-differences estimates on ln Emissions across pollutants from Section 3 and their model counterparts are plotted. The results are calculated from Equation (2). I apply the Frisch–Waugh–Lovell weights implied by the regression's fixed effects to simulated cell-level log emissions changes. The last two bars show the model-implied changes in regulated and unregulated cells; the "Model" bar equals the first minus the second.
Sources: Ministry of Environment.

Table 7: Sources of emissions reductions in regulated districts

Pollutant	Equipment	Fuel intensity and mix	Scale and reallocation	Total
NO_x	−17.8	−2.3	−13.3	−33.3
SO_x	−11.5	−0.2	−20.5	−32.1
TSP	−27.2	−0.5	−20.4	−48.0

Notes: Entries are percentage-point contributions relative to each pollutant's total baseline emissions in regulated districts, including all manufacturing and power plants regardless of plant-level eligibility. Manufacturing emissions are decomposed into equipment intensity, pre-abatement emissions per unit of output, and plant-level physical output. The second component changes through fuel intensity and composition; the output-linked emissions coefficient is fixed. Contributions average over all six orders of changing these factors. Power emissions are decomposed into the equipment factor, net of its generation loss, and net generation, averaging over both orders. Scale and reallocation includes changes in output across plants, not only changes in aggregate output. This is an accounting decomposition of simulated emissions changes, not a set of separate equilibrium counterfactuals. Totals may differ from sums of displayed contributions due to rounding.

districts account for the remaining 19.1%.[26] However, geographic pollution shifts in power generation are weak. Both the model and the empirical estimates show larger relative emissions declines in manufacturing than in power generation (Section 3).

To understand the sources of emissions reductions, I present a Shapley decomposition of emissions reductions in regulated districts in Table 7. Equipment use is the largest contributor to NO_x and TSP reductions, accounting for 17.8 and 27.2 percentage points, respectively, relative to baseline emissions. Changes in plant-level output and its allocation contribute most to the SO_x reduction, accounting for 20.5 percentage points. Changes in fuel intensity and composition make smaller contributions for all three pollutants.

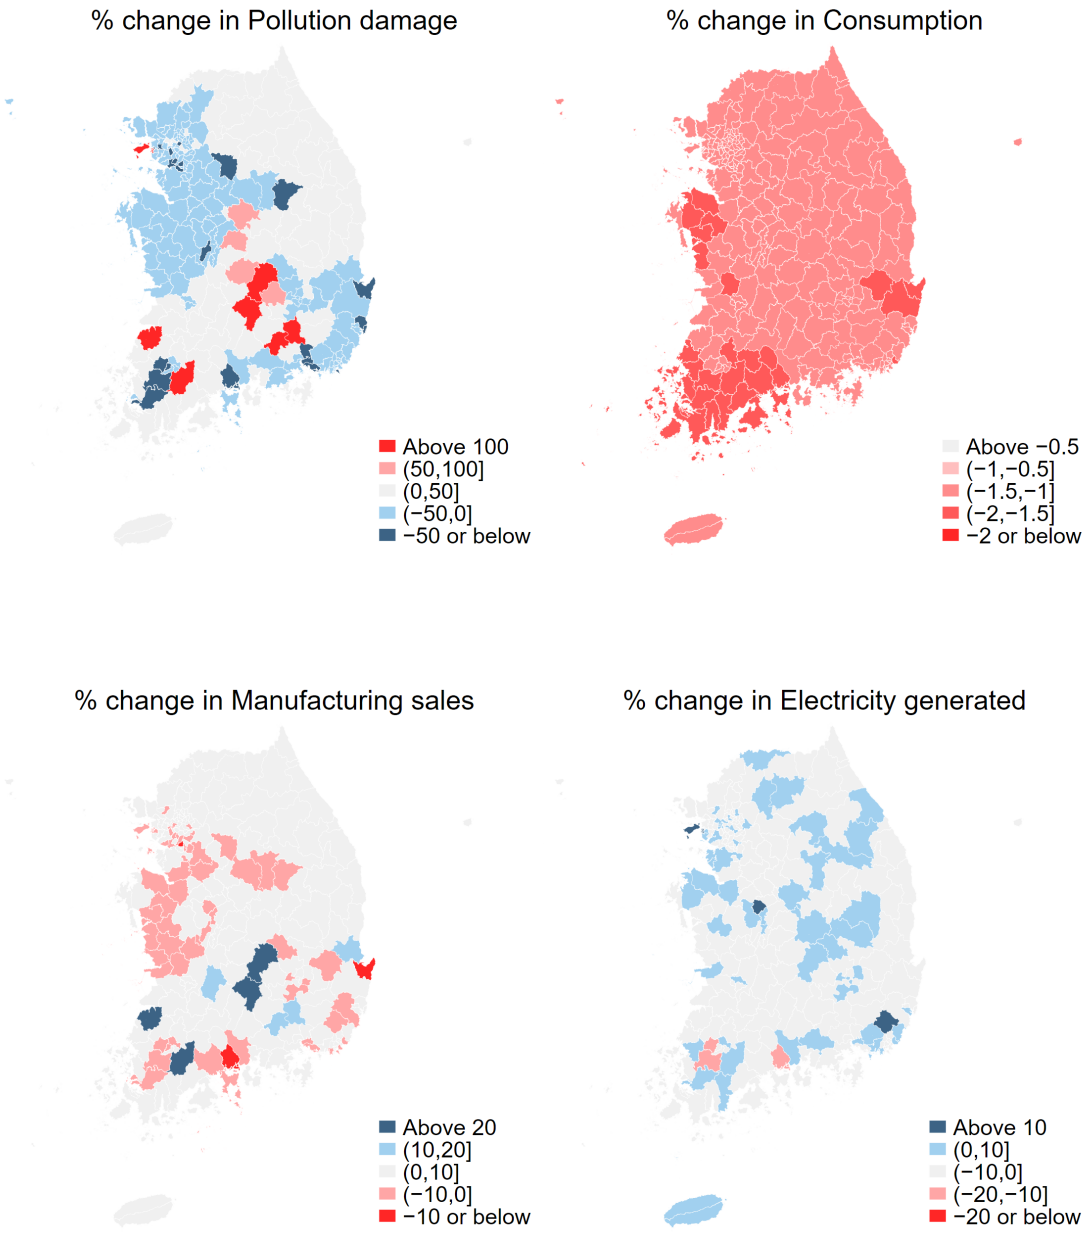
Taken together, these findings motivate tracing how the policy changes the spatial distribution of emissions and economic activity.

6.1.2 Consumption, Pollution Damage, and Distributional Consequences

Figure 7 shows district-level changes in pollution damage and consumption in the upper panels and in manufacturing sales and electricity generation in the lower panels. The spatial effects differ sharply between regulated and unregulated districts. Manufacturing sales fall by 0.4% in regulated districts and rise by 4.5% outside them, relative to each group's baseline sales. Aggregate manufacturing sales remain essentially unchanged under fixed district–sector expenditure, so these changes reflect spatial reallocation. Similarly,

[26] In a related decomposition for the U.S. Clean Air Act, Curtis et al. (2026) attribute about 40% of the relative employment decline to expansion in unregulated locations.

Figure 7: Distributional consequences across space



Notes: Changes are measured relative to their own unregulated-equilibrium levels. I use sales rather than physical output to aggregate manufacturing activity across sectors.

Table 8: Counterfactual policy experiments

Policy	District group	Consumption change (%)	Pollution damage change (%)
Actual policy	All	−1.42	−37.8
	Regulated	−1.43	−39.6
	Unregulated	−1.36	27.9
Uniform pricing	All	−0.82	−32.1
	Regulated	−0.82	−31.5
	Unregulated	−0.85	−55.3
Damage-proportional pricing	All	−0.98	−41.8
	Regulated	−0.98	−41.7
	Unregulated	−0.95	−45.0

Notes: Consumption and pollution damage are reported as percentage changes relative to their respective baseline totals. Regulated and unregulated districts retain their actual policy classification. "Uniform pricing" denotes nationwide eligibility under the same calibrated thresholds, with common emissions prices chosen to target actual-policy national emissions totals. "Damage-proportional pricing" denotes nationwide eligibility with district prices proportional to baseline marginal damages and pollutant-specific price levels chosen to target the actual policy's national emissions totals. See Table A.28 for area-specific changes in the actual policy case.

national electricity generation is essentially unchanged, while generation shifts across districts. The resulting spatial distribution of emissions matters for aggregate pollution damage because pollution damages depend on pollutant-specific damage coefficients and local population exposure.

The first three rows of Table 8 report the aggregate and spatial effects of the actual policy. Pollution damage falls by 39.6% in regulated districts but rises by 27.9% outside them. Consumption falls by 1.43% and 1.36%, respectively. Aggregate consumption falls by 1.42%, which reflects higher consumption prices.[27] Aggregate pollution damage falls by 37.8% because damage reductions in regulated districts outweigh increases elsewhere. For local air pollutants, the location of emissions matters: both reductions in high-damage districts and reallocation toward lower-damage districts can reduce aggregate damage. Climate damages from greenhouse gases, by contrast, do not depend on the location of emissions in the same way.

[27]Related models also imply consumption costs through higher prices. Curtis et al. (2026) report an expenditure-weighted price increase of about 2% in regulated U.S. manufacturing industries under the U.S. Clean Air Act, while Chen et al. (2025) report a price-index increase of about 3.5% in their model of China's energy-intensive manufacturing industries under China's Top 1,000 program.

6.2 Alternative Policy Designs

The costs and benefits of regulation depend on both the cost of reducing emissions and where those reductions occur. I use the spatial model to consider two nationwide alternatives: a common emissions pricing rule and a damage-proportional pricing rule, both targeting the actual policy's national emissions totals.

First, to assess a policy response to geographic pollution shifts, I extend eligibility nationwide, applying the same calibrated pollutant-specific thresholds to baseline emissions in every district. Qualifying plants in previously unregulated districts therefore face regulation as well. I choose three national emissions prices, one per pollutant, to target the actual policy's national emissions totals. The price adjustments differ across pollutants.

Second, I retain nationwide eligibility but set district emissions prices in proportion to marginal damage weights, which combine pollutant-specific social costs with local population exposure. Holding these relative prices fixed, I choose three non-negative pollutant-specific base prices to match the actual policy's national emissions totals. Unlike the actual policy's broad regional pricing structure, this rule assigns higher prices within each pollutant to districts with greater population-exposure weights.

The first experiment is reported in the three rows in the middle of Table 8. Nationwide eligibility limits incentives to shift production toward geographically exempt plants and broadens the set of plants contributing to the same national emissions targets. Under uniform nationwide prices, aggregate pollution damage falls by 32.1% at the cost of a 0.82% decline in consumption, compared with 37.8% and 1.42%, respectively, under the actual policy.

The smaller aggregate damage reduction reflects a change in where emissions fall. Uniform prices do not give greater weight to reductions in high-damage locations. Pollution damage in originally regulated districts falls by 31.5%, compared with 39.6% under the actual policy. In previously unregulated districts, damage instead falls by 55.3%, rather than rising by 27.9%, while consumption falls by 0.85%. Because baseline damage is concentrated in originally regulated districts, the smaller reductions there still outweigh the additional benefits elsewhere. Attenuating geographic pollution shifts therefore need not deliver a larger reduction in aggregate pollution damage when national emissions totals are held fixed.

The second experiment is reported in the last three rows of Table 8. Aggregate pollution damage falls by 41.8% at the cost of a 0.98% decline in consumption. This alternative achieves a larger damage reduction and a smaller consumption loss than the actual policy. Compared with uniform pricing, however, its greater environmental benefits come with a larger consumption loss. Both nationwide alternatives extend eligibility to the same

previously excluded districts and target the same national emissions totals. Their difference in aggregate damage therefore reflects the spatial price rule: damage-proportional pricing directs stronger relative incentives toward locations with higher marginal damage.

7 Conclusion

Three findings emerge from the analysis. First, manufacturing emissions fall by 32.9–36.9% across pollutants in regulated districts relative to the same industries in unregulated districts. These reductions coexist with production reallocation within manufacturing firms and across districts, as well as potential emissions shifts to power generation. Second, the spatial model shows that geographic pollution shifts account for 19.1% of the manufacturing emissions difference-in-differences estimate. Third, the actual policy reduces aggregate pollution damage by 37.8% at the cost of a 1.42% decline in consumption. Pollution damage falls in regulated districts but rises outside them, while consumption declines in both groups. The policy's benefits therefore depend on where emissions change, not only on the national reduction.

At the same pollutant-specific national emissions totals, nationwide coverage attenuates geographic pollution shifts, but uniform prices do not prioritize reductions in high-damage locations. Relative to the actual policy, this design entails smaller consumption losses but leaves emissions reductions less concentrated in these locations, yielding smaller aggregate damage reductions. Damage-proportional pricing adds spatially differentiated incentives to nationwide coverage, directing stronger incentives toward locations with higher marginal damage. It achieves greater damage reductions than either the actual policy or uniform pricing, with consumption costs below those of the actual policy but above those of uniform pricing. These results show that the benefits of regulating local air pollution depend on both the cost of reducing emissions and the spatial distribution of the resulting damages.

References

Abadie, Alberto and Jann Spiess (2022). Robust Post-Matching Inference. *Journal of the American Statistical Association* 117.538, pp. 983–995.

Arkolakis, Costas, Fabian Eckert, and Rowan Shi (2023). Combinatorial Discrete Choice: Theory and Application to Multinational Production. NBER Working Paper 31877. National Bureau of Economic Research. URL: https://www.nber.org/papers/w31877.

Arkolakis, Costas, Cheolhwan Kim, and Emmanuel Murray Leclair (2025). Deep Industrial Decarbonization: Theory and Evidence from South Korea. Working paper.

Arkolakis, Costas and Conor Walsh (2023). Clean Growth. NBER Working Paper 31615. National Bureau of Economic Research. URL: https://www.nber.org/papers/w31615.

Blundell, Wesley, Gautam Gowrisankaran, and Ashley Langer (2020). Escalation of Scrutiny: The Gains from Dynamic Enforcement of Environmental Regulations. *American Economic Review* 110.8, pp. 2558–2585.

Broda, Christian and David E. Weinstein (2006). Globalization and the Gains from Variety. *The Quarterly Journal of Economics* 121.2, pp. 541–585.

Chen, Qiaoyi, Zhao Chen, Zhikuo Liu, Juan Carlos Suárez Serrato, and Daniel Yi Xu (2025). Regulating Conglomerates: Evidence from an Energy Conservation Program in China. *American Economic Review* 115.2, pp. 408–447.

Christensen, Laurits R. and William H. Greene (1976). Economies of Scale in U.S. Electric Power Generation. *Journal of Political Economy* 84.4, pp. 655–676.

Coase, Ronald H. (1960). The Problem of Social Cost. *The Journal of Law and Economics* 3, pp. 1–44.

Colmer, Jonathan, Ralf Martin, Mirabelle Muûls, and Ulrich J. Wagner (2025). Does Pricing Carbon Mitigate Climate Change? Firm-Level Evidence from the European Union Emissions Trading System. *The Review of Economic Studies* 92.3, pp. 1625–1660.

Copeland, Brian R. and M. Scott Taylor (2004). Trade, Growth, and the Environment. *Journal of Economic Literature* 42.1, pp. 7–71.

Cruz, José-Luis and Esteban Rossi-Hansberg (2025). Local Carbon Policy. In: *Advances in Economics and Econometrics: Twelfth World Congress*. Ed. by Victor Chernozhukov, Johannes Hörner, Eliana La Ferrara, and Iván Werning. Vol. 2. Cambridge University Press, pp. 292–332.

Cui, Jingbo, Chunhua Wang, Junjie Zhang, and Yang Zheng (2021). The Effectiveness of China's Regional Carbon Market Pilots in Reducing Firm Emissions. *Proceedings of the National Academy of Sciences* 118.52, e2109912118.

Curtis, E. Mark and Jonathan M. Lee (2019). When Do Environmental Regulations Backfire? Onsite Industrial Electricity Generation, Energy Efficiency and Policy Instruments. *Journal of Environmental Economics and Management* 96, pp. 174–194.

Curtis, E. Mark, Weiting Miao, Felix Samy Soliman, Juan Carlos Suárez Serrato, and Daniel Yi Xu (2026). The Geographical Leakage of Environmental Regulation: Evidence from the Clean Air Act. Working paper. URL: https://fsamysoliman.github.io/Curtis-Miao-Soliman-Serrato-Xu-compressed.pdf.

De Loecker, Jan (2011). Product Differentiation, Multiproduct Firms, and Estimating the Impact of Trade Liberalization on Productivity. *Econometrica* 79.5, pp. 1407–1451.

De Simone, Rebecca, S. Lakshmi Naaraayanan, and Kunal Sachdeva (2024). Opening the Brown Box: Production Responses to Environmental Regulation. SSRN Working Paper 4707581. Social Science Research Network. URL: https://ssrn.com/abstract=4707581.

Dechezleprêtre, Antoine, Daniel Nachtigall, and Frank Venmans (2023). The Joint Impact of the European Union Emissions Trading System on Carbon Emissions and Economic Performance. *Journal of Environmental Economics and Management* 118, 102758.

Dixit, Avinash K. and Joseph E. Stiglitz (1977). Monopolistic Competition and Optimum Product Diversity. *American Economic Review* 67.3, pp. 297–308.

Doraszelski, Ulrich and Jordi Jaumandreu (2013). R&D and Productivity: Estimating Endogenous Productivity. *The Review of Economic Studies* 80.4, pp. 1338–1383.

— (2018). Measuring the Bias of Technological Change. *Journal of Political Economy* 126.3, pp. 1027–1084.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan (2013). Truth-Telling by Third-Party Auditors and the Response of Polluting Firms: Experimental Evidence from India. *The Quarterly Journal of Economics* 128.4, pp. 1499–1545.

Farrokhi, Farid and Ahmad Lashkaripour (2025). Can Trade Policy Mitigate Climate Change? *Econometrica* 93.5, pp. 1561–1599.

Fowlie, Meredith (2009). Incomplete Environmental Regulation, Imperfect Competition, and Emissions Leakage. *American Economic Journal: Economic Policy* 1.2, pp. 72–112.

— (2010). Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement. *American Economic Review* 100.3, pp. 837–869.

Fowlie, Meredith, Stephen P. Holland, and Erin T. Mansur (2012). What Do Emissions Markets Deliver and to Whom? Evidence from Southern California's NO_x Trading Program. *American Economic Review* 102.2, pp. 965–993.

Fowlie, Meredith and Jeffrey M. Perloff (2013). Distributing Pollution Rights in Cap-and-Trade Programs: Are Outcomes Independent of Allocation? *The Review of Economics and Statistics* 95.5, pp. 1640–1652.

Fowlie, Meredith, Mar Reguant, and Stephen P. Ryan (2016). Market-Based Emissions Regulation and Industry Dynamics. *Journal of Political Economy* 124.1, pp. 249–302.

Gillingham, Kenneth and Pei Huang (2024). Racial Disparities in the Health Effects from Air Pollution: Evidence from Ports. *Journal of Human Resources*.

Greenstone, Michael and Christa Hasenkopf (2023). Air Quality Life Index: 2023 Annual Update. Report. Energy Policy Institute at the University of Chicago. URL: `https://aqli.epic.uchicago.edu`.

Greenstone, Michael, Rohini Pande, Nicholas Ryan, and Anant Sudarshan (2025). Can Pollution Markets Work in Developing Countries? Experimental Evidence from India. *The Quarterly Journal of Economics* 140.2, pp. 1003–1060.

Grieco, Paul L. E., Shengyu Li, and Hongsong Zhang (2016). Production Function Estimation with Unobserved Input Price Dispersion. *International Economic Review* 57.2, pp. 665–690.

Hahn, Robert W. and Robert N. Stavins (2011). The Effect of Allowance Allocations on Cap-and-Trade System Performance. *The Journal of Law and Economics* 54.S4, S267–S294.

Heckman, James J., Hidehiko Ichimura, and Petra E. Todd (1997). Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme. *The Review of Economic Studies* 64.4, pp. 605–654.

— (1998). Matching as an Econometric Evaluation Estimator. *The Review of Economic Studies* 65.2, pp. 261–294.

Heo, Seonmin Will, Koichiro Ito, and Rao Kotamarthi (2025). International Spillover Effects of Air Pollution: Evidence from Mortality and Health Data. *The Review of Economics and Statistics*.

Heshmati, Almas, Subal C. Kumbhakar, and Kai Sun (2014). Estimation of Productivity in Korean Electric Power Plants: A Semiparametric Smooth Coefficient Model. *Energy Economics* 45, pp. 491–500.

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* 4.1, pp. 38–77.

Im, Jiwon (2023). An Empirical Study on the Effects of the Cap-and-Trade System on Air Pollutants in the Seoul Metropolitan Area Industrial Settings. MA thesis. Seoul, Korea: Seoul National University.

Intergovernmental Panel on Climate Change (2022). Climate Change 2022: Mitigation of Climate Change. Working Group III Contribution to the Sixth Assessment Report. Chapter 2. Intergovernmental Panel on Climate Change. URL: https://www.ipcc.ch/report/ar6/wg3/.

International Energy Agency (2021). Air Quality and Climate Policy Integration in India: Frameworks to Deliver Co-Benefits. Report. International Energy Agency.

Jia, Wentao, Chunbo Ma, Xiao Wang, and Rui Xie (2025). Intra-Firm Pollution Leakage and Redistribution of Pollution Exposure: Evidence from Coal-Regulated Plants in China. *Journal of Environmental Economics and Management* 133, 103205.

Jiang, Jingjing, Bin Ye, Zhenzhong Zeng, Xin Yang, Zhuoluo Sun, Shuai Shao, Kuishuang Feng, and Xiujie Tan (2024). Carbon Abatement and Leakage in China's Regional Carbon Emission Trading. *Environmental Science & Technology* 58.40, pp. 17661–17673.

Korea Power Exchange (2020). A Study on Current Trends in the Levelized Cost of Electricity by Generation Source. Report. In Korean. Naju, Korea: Korea Power Exchange.

— (2022). Domestic Trends: 2022 Survey of Commercial Self-Generation Companies. Report. In Korean. Naju, Korea: Korea Power Exchange.

Kortum, Samuel and David A. Weisbach (2024). Optimal Unilateral Carbon Policy. Cowles Foundation Discussion Paper 2311R1. Cowles Foundation for Research in Economics, Yale University. URL: https://cowles.yale.edu/research/cfdp-2311r1-optimal-unilateral-carbon-policy.

Lee, Hye-kyung (2019). The Industrial Air Pollution Total Emissions Management System: Current Status and Improvements. Legislative and Policy Report 36. In Korean. Seoul, Korea: National Assembly Research Service. URL: https://www.nars.go.kr/report/view.do?brdSeq=26792&cmsCode=CM0156.

Levinson, Arik and M. Scott Taylor (2008). Unmasking the Pollution Haven Effect. *International Economic Review* 49.1, pp. 223–254.

MacKinnon, James G. and Matthew D. Webb (2020). Randomization Inference for Difference-in-Differences with Few Treated Clusters. *Journal of Econometrics* 218.2, pp. 435–450.

Martin, Ralf, Mirabelle Muûls, Laure B. de Preux, and Ulrich J. Wagner (2014). Industry Compensation under Relocation Risk: A Firm-Level Analysis of the EU Emissions Trading Scheme. *American Economic Review* 104.8, pp. 2482–2508.

Melitz, Marc J. (2003). The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica* 71.6, pp. 1695–1725.

Muller, Nicholas Z. and Robert Mendelsohn (2009). Efficient Pollution Regulation: Getting the Prices Right. *American Economic Review* 99.5, pp. 1714–1739.

Murray Leclair, Emmanuel (2024). Balancing Production and Carbon Emissions with Fuel Substitution. Working paper.

Naegele, Helene and Aleksandar Zaklan (2019). Does the EU ETS Cause Carbon Leakage in European Manufacturing? *Journal of Environmental Economics and Management* 93, pp. 125–147.

Oh, Dong-hyun and Yong-Gil Lee (2016). Productivity Decomposition and Economies of Scale of Korean Fossil-Fuel Power Generation Companies: 2001–2012. *Energy* 100, pp. 1–9.

Parry, Ian W. H., Dirk Heine, Eliza Lis, and Shanjun Li (2014). *Getting Energy Prices Right: From Principle to Practice*. Washington, DC: International Monetary Fund.

Ryan, Nicholas (2018). Energy Productivity and Energy Demand: Experimental Evidence from Indian Manufacturing Plants. NBER Working Paper 24619. National Bureau of Economic Research.

Shapiro, Joseph S. and Reed Walker (2018). Why Is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade. *American Economic Review* 108.12, pp. 3814–3854.

— (2025). Is Air Pollution Regulation Too Lenient? Evidence from US Offset Markets. *American Economic Review* 115.9, pp. 3058–3080.

Tanaka, Shinsuke, Kensuke Teshima, and Eric Verhoogen (2022). North-South Displacement Effects of Environmental Regulation: The Case of Battery Recycling. *American Economic Review: Insights* 4.3, pp. 271–288.

Tang, Ling, Min Jia, Junai Yang, Ling Li, Xin Bo, and Zhifu Mi (2023). Chinese Industrial Air Pollution Emissions Based on the Continuous Emission Monitoring Systems Network. *Scientific Data* 10, 153.

Tintelnot, Felix (2017). Global Production with Export Platforms. *The Quarterly Journal of Economics* 132.1, pp. 157–209.

von Graevenitz, Kathrine and Elisa Rottner (2026). Climate Policies and Electricity Prices: Why Abate When You Can Generate? *American Economic Journal: Economic Policy*. Forthcoming.

Weisbach, David A., Samuel Kortum, Michael Wang, and Yujia Yao (2023). Trade, Leakage, and the Design of a Carbon Tax. *Environmental and Energy Policy and the Economy* 4, pp. 43–90.

Xiang, Wei (2023). Clean Growth and Environmental Policies in the Global Economy. SSRN Working Paper 4619717. Social Science Research Network. URL: https://ssrn.com/abstract=4619717.

Zhang, Bing and Daxuan Zhao (2023). Emission Leakage and the Effectiveness of Regional Environmental Regulation in China. *Journal of Environmental Economics and Management* 121, 102869.

A Institutional Background

A.1 Policy Details

A.1.1 Rules on Caps

Initial caps are allocated based on past emissions volumes. To protect early adopters, establishments that have already invested in abatement technology are assessed against their five-year average emissions.[28] Subsequent permit allocations are benchmarked to "Best Available Technology" standards, which are established in consultation with industry and reviewed every three years. Importantly, each area's reduction goal is based on its baseline emissions and population, which is the source of the policy variation by district.

A.1.2 Trading

The policy was implemented at the area level (a set of neighboring districts), not at the district level, to avoid market frictions from small market sizes. In the pollution markets, an establishment or a firm can buy or sell its permits on a website designed and controlled by the Korean Ministry of Environment by posting the prices at which it is willing to buy or sell permits (see Stack and Sky permit trading platform). The markets are decentralized, so buyers and sellers write bilateral contracts, which have to be certified by the ministry. Although the system collects bids, it is closer to bilateral trading than to an auction because both parties negotiate the terms.[29] Brokerage is not allowed. Banking is allowed for one year while borrowing is prohibited. The static model abstracts from the intertemporal incentives created by one-year banking.

Permit demand was concentrated outside the energy sector, mainly among manufacturing and waste/other facilities, whereas supply was concentrated among power plants, in terms of both frequencies and quantities. During 2020–2023, power plants accounted for 45% of total intended sales but only 23% of total intended purchases. This asymmetry suggests that power plants were important prospective permit suppliers, while non-energy firms were the main demand side of the market. This pattern is consistent with power plants entering the policy period with favorable permit balances, potentially reflecting their initial allocations as early adopters.

Lastly, note that the NO_x and SO_x prices are of a similar order of magnitude to permit prices in major emissions-trading systems (ETS), such as the EU ETS, California, and China's national ETS. However, these are not directly comparable because they cap

[28] I abstract from the allocation of permits for two reasons. First, under standard conditions, a Coasian argument implies that the equilibrium allocation of abatement across firms is cost-effective and independent of the initial distribution of permits, although the initial allocation may still affect the distribution of compliance burdens (Coase 1960). Second, I do not have a credible empirical strategy to identify the causal effect of permit allocation itself. For related discussions, see Hahn and Stavins (2011) and Fowlie and Perloff (2013).

[29] Because many bids report negotiated or zero prices, explicit bid prices are highly selected. I therefore do not use bid prices to estimate marginal abatement costs. Transactions without consideration are rare, but arise between plants within the same ownership.

different pollutants (NO_x and SO_x vs. CO_2), with different damages, sector coverage, monitoring rules, and compliance designs. A more directly comparable benchmark is the U.S. Cross-State Air Pollution Rule (CSAPR) market for NO_x and SO_2. Korea's NO_x and SO_x prices are above typical U.S. annual CSAPR permit prices but below U.S. ozone-season NO_x permit prices, which can be much higher. TSP has rarely been regulated, so it is hard to find a counterpart. One exception can be found in Greenstone et al. (2025).

A.1.3 Measurement, Reporting, and Verification

Annual emissions are measured over the calendar year, from January through December. Measurement, Reporting, and Verification (MRV), in principle, rely on sensors at establishment outlets. In practice, emissions are measured continuously by the TMS where available; otherwise, plants report self-measurement results and operating data through the SEMS, including measurements conducted by certified testing agencies. Before using the SEMS submissions to calculate emissions, the National Air Emission Inventory and Research Center verifies them by cross-checking them against other administrative data. Regulated establishments must surrender permits annually to match realized emissions. Excess emissions trigger fines, reductions in future permit allocations, and potentially shutdown orders, with penalties increasing in the severity of the violation.

A.1.4 Other Details

The program itself includes permitting requirements for qualifying plant installations and modifications. It applies to new or modified emissions sources rather than mandating reductions from unchanged existing sources. These permits require best available control facilities. Previously applicable in the Seoul Capital Area, this requirement expanded with the program in 2020.

Although the policy primarily targets large stationary sources, the government introduced financial support to ease compliance costs for small and medium-sized firms subject to the expanded regulatory regime. For example, it subsidized up to 90% of the installation costs for emission-control facilities and up to 80% of the installation and maintenance costs for the TMS.

A.1.5 Policy Timing

The initial phase was from 2007 to 2019 in the Seoul Capital Area only. Then, the program expanded to the other three areas in 2020. This phase was from 2020 to 2024. A new phase began in 2025 with several changes, including a borrowing system and an external reduction activity recognition system introduced on August 17, 2024. The empirical sample covers 2016–2023 and therefore excludes the later rule changes and focuses on the second phase.

The difference-in-differences design assumes no material anticipatory responses during the pre-policy comparison period. For the Seoul Capital Area where the regulation existed earlier, I treat 2020 as a policy intensity update and present robustness checks excluding this area. The enabling statute was enacted in April 2019. The proposed geographic

coverage and area-specific public briefings were announced in November 2019. The policy took effect in April 2020, when regional caps were established. Plant-level allocations for 2020–2024 were announced in October 2020. This sequence suggests that plants faced uncertainty about policy details before implementation.

Eligibility relied on absolute, recent-year criteria for sizable sources in the new control zones, which are difficult to fine-tune just before implementation. Initial allocations in 2020 were based on 2019 emissions, with adjustments using five-year average emissions to avoid penalizing early abatement. These policies target large polluters in populated and polluted districts, prioritizing public health rather than climate mitigation. Eligibility based on historical emissions limited establishments' ability to avoid regulation through last-minute adjustments, even if they anticipated the expansion.

A.2 Electricity Market

Korea's electricity market has separate wholesale and retail segments: a wholesale market operates between power plants and the government, while a retail market connects the government to final consumers such as households and manufacturing plants. The wholesale market operates through two commonly used mechanisms: system marginal price (SMP) setting and merit-order dispatch (MOD). Lower-cost plants, such as nuclear or renewable generators, enter before fossil-fuel generators (MOD), and the electricity price is set by the marginal cost of the final generation increment dispatched (SMP).

Yet given limited clean capacity, short-run demand surges likely call on dirtier power plants because their marginal costs are higher and clean power plants are already dispatched. Clean power generation accounted for 28% of total capacity in 2018. Given merit-order dispatch and the uneven location of generation capacity, additional demand can be met by fossil-fuel generators in other districts. The resulting emissions shifts depend on the amount and location of available clean capacity.

Because eligibility is based on pre-policy emissions, all initially regulated power plants had already generated electricity before the policy; a plant entering the merit order for the first time after implementation would therefore not belong to the initially regulated group. By raising marginal costs, regulation can move a previously dispatched plant lower in the dispatch ranking or leave it undispatched. If displaced, its output and emissions fall; if it remains dispatched, higher manufacturing electricity demand can maintain or increase its generation and offset the direct compliance effect. Provided that aggregate electricity demand does not fall, higher generation costs imply that the electricity price remains unchanged or rises. The absence of a large sales decline among regulated power plants, together with increased electricity use among regulated manufacturers, is consistent with the latter case (Section 3).

The mainland wholesale market uses a common SMP, while Jeju has a separate SMP. However, all the other features are standard as in other countries (see, e.g., Arkolakis and Walsh 2023). The quantitative model inherits these market mechanisms. Due to well-established grids, transmission and distribution are a minor issue in South Korea. Zonal pricing is under discussion but not yet implemented. A key difference is that the model integrates the wholesale and retail markets into one. Since the frequency is yearly and the paper focuses on short- or medium-run effects of the policy, this simplification

abstracts from subsidies to electricity consumers and generators. However, the relative patterns are matched in the model without such features, and these other policies are beyond the scope of this paper.

B Data Details

This section provides additional details on the data construction. I group mining with manufacturing because of their similar resource-based production processes and the frequent switching of plants between the two categories. Then, I clean and use each dataset as follows.

B.1 Plant-Level Data

The SBR covers the universe of firms and establishments. I apply five data-cleaning procedures. First, I fill in missing location codes where possible. For example, when district codes are missing, I infer them from more granular location codes. Second, I harmonize industry and district codes using the 2023 classifications. For KSIC codes, I set pre-2017 codes to missing, because the 10th revision of the KSIC was implemented in 2017. Then, I fill in missing industry codes using each plant's available codes from 2017, 2018, and 2019. Third, I use firm and plant IDs, districts, KSIC codes, latitudes and longitudes, and histories to identify the same plants within firms in specific districts. Fourth, I merge with the MMS and the equipment use dataset using plant IDs and latitudes and longitudes. The MMS provides detailed information for all mining and manufacturing plants with 10 or more employees and a representative sample of smaller plants. Lastly, I use the registry of regulated plants containing information about targeted pollutants, locations, sectors, and plant names.

Emissions are also observed in real time at the plant level with the tele-monitoring sensors, but primarily for large polluters. This sample covers too few untreated plants to form a plausible control group. While the fuel use survey has informed industrial environmental policies, it has never been used directly for policy implementation. Misreporting is less likely because the Ministry of Environment cross-checks records with other relevant data (e.g., real-time monitored data) to scrutinize the survey responses. This issue is more likely to be prominent in developing countries where cross-checks are challenging. For such an issue, see Duflo et al. (2013).[30] Electricity use includes electricity generated by plants themselves. But it is more likely to be purchased electricity because on-site generation is

[30]I cannot comprehensively assess bunching or responses to the threat of regulation because doing so would require plant-level emissions data for the full universe of regulated and unregulated plants. From the list of regulated plants, I find that the regulatory status of plants barely changes. Still, unregulated plants within regulated districts might bunch below the thresholds. Using a subset of the universe, I plot the emissions by pollutant. The histograms show no clear bunching immediately below the pollutant-specific cutoffs (see Figure A.6). For robustness, I check plant-level sales and fossil-fuel use for plants in regulated districts and find no conclusive evidence. Duflo et al. (2013) note that emissions variability makes it difficult to target a regulatory standard precisely and document bunching in reported pollution concentrations caused by misreporting.

rare: non-power plants account for only about 3.8% of total generation capacity, and only 32 manufacturing firms generate electricity on-site (Korea Power Exchange 2022).

The raw policy registry contains 1,319 manufacturing and power plants before linkage to administrative and census data or the application of panel-balance restrictions. Fifteen listed plants had closed before the sample period but remained on the policy list. The 1,093 plants reported in the main text are those identified in the merged data, rather than the total number in the raw registry. Before linkage to the census, plant names and addresses serve as plant identifiers in the policy registry.

Firms are defined at the corporate-entity level. For non-corporate entities, I define firms by common ownership. I do not define firms at the chaebol level because doing so aggregates distinct corporate entities into broad business groups. To account for mergers, acquisitions, and splits, I supplement the census with records from the Data Analysis, Retrieval and Transfer System of the Korean Financial Supervisory Service. Business leases are recorded from the lessor's perspective, consistent with the underlying data.

B.2 District–Industry-Level Data

The CAPSS combines tele-monitored emissions for covered large sources with estimates based on reported fuel use, production processes, emission factors, and abatement equipment for other establishments.[31] For combustion emissions from small manufacturing plants, the CAPSS estimates emissions using fuel statistics net of consumption recorded in SEMS and allocates these emissions across districts. The CAPSS provides nationwide emissions estimates used to inform environmental policy. Continuous emissions monitoring through the TMS is required for designated stacks at large establishments, subject to applicable exemptions. Other sources are covered through periodic establishment surveys and emissions estimation. This difference in reporting based on emissions volume reflects the fact that large polluters account for most emissions. For example, small polluters were surveyed every four years before 2021 and annually thereafter. Indeed, although the sensors were installed only on a few large plants (less than 1% of the population), from 2016 to 2022, 63%, 45%, and 22% of total NO_x, SO_x, and TSP emissions from manufacturing and power plants were directly tele-monitored when I compare tele-monitored emissions and total emissions from the CAPSS.

I harmonize district codes in the CAPSS and census data using the 2023 district classifications. The CAPSS emissions are classified by emission source while the SBR is categorized by the classification of the product with the highest sales. Most emissions-source categories can be mapped to manufacturing KSIC codes at the two- to five-digit level and to five-digit codes in power generation. To avoid imputing emissions across classifications, I retain the original categories where they align and otherwise aggregate them to a common level.

I merge the CAPSS emissions data with the SBR and then the pollution-market-level panel that has information about emissions, permit allocations, and permit prices over time. The match is complete.

[31] This information is available on the Ministry of Environment's website: National Air Emission Inventory and Research Center.

B.3 Emission Factors

Pollutant- and source-specific emission factors, developed by the Ministry of Environment jointly with the Korean National Institute of Environmental Research, are available for emissions from production processes and fossil-fuel use. Due to the volume of the tables, they are omitted, but they are available upon request.

B.4 Social Costs of Local Air Pollutants

Although the literature typically estimates these values using integrated assessment models, I construct an initial measure of pollution damage using pollutant-specific unit damage values reported by Korea Power Exchange (2020), which adapts estimates from Parry et al. (2014). These estimates combine information on Korean power-plant locations and population exposure with health-risk valuations: particulate matter causes direct exposure, whereas NO_x and SO_x generate damages chiefly through secondary particulate formation. I convert the reported values to U.S. dollars using the KRW–USD exchange rate from the Federal Reserve Bank of St. Louis: 34.3 USD/kg for NO_x, 47.5 USD/kg for SO_x, and 62.1 USD/kg for particulate matter. Because the emissions data report TSP rather than size-specific particulate matter, I apply the particulate-matter coefficient as a proxy.

C Details about Difference-in-Differences

Table A.3 shows the sectoral distribution of regulated plants. As expected, they are concentrated in heavy manufacturing and fossil-fuel power generation sectors. Tables A.4 and A.5 present summary statistics at both levels. At the district–industry level, regulated manufacturing has higher emissions than unregulated manufacturing, while regulated power generation has lower emissions. The location of fossil-fuel power plants away from populated districts means that manufacturing electrification can shift emissions across districts. At the plant level, regulated plants are larger and exhibit more abatement efforts.

C.1 Matching

At the plant level, I construct matched difference-in-differences comparisons following Heckman et al. (1997) and Heckman et al. (1998). The goal of matching is to compare regulated plants with observationally similar unregulated plants, while recognizing that the resulting estimates are relative to the matched unregulated plants and may reflect spillovers to those plants.

Matching proceeds in three steps. First, I exclude from the control group plants that were not subject to any environmental regulations before the cap-and-trade expansion, since they may face different compliance requirements or operate under different production and emissions technologies. For the comparison between unregulated affiliated and unaffiliated manufacturing plants, I apply the same restriction. I also report robustness

checks about this trimming rule. Second, I use strongly balanced samples with non-missing outcome variables.

Third, I exactly match plants by three-digit KSIC in 2017 and, for manufacturing plants, by policy area. I do not impose exact area matching for power plants because treated power plants are sparse. The year 2017 is chosen because this is when a new KSIC system was introduced. Within these exact-match cells, I use one-to-one nearest-neighbor Mahalanobis distance matching without replacement. The distance variables are pre-period averages of ln Electricity use and ln Sales for manufacturing plants, and pre-period average ln Sales for power plants. I use the pre-period average to reduce sensitivity to idiosyncratic annual shocks. Thus, controls for regulated manufacturing plants are unregulated unaffiliated plants below the regulatory threshold in regulated districts, controls for unregulated affiliates are unaffiliated plants in the same policy area, and controls in power generation are either unregulated plants below the regulatory threshold in regulated districts or plants in unregulated districts.

The final matched samples are smaller than the initial samples, but matching improves balance in pre-policy sales and electricity use.[32] Matching worsens balance in pre-policy equipment updating: matched unregulated plants have higher initial update rates, whereas regulated plants exhibit a larger increase in updating after the policy.

D Robustness Checks

This section presents robustness checks on the reduced-form analysis. As above, I use two units of analysis: district–industry cells and plants. In general, the main qualitative patterns are preserved across the robustness checks. I start by addressing two factors that may confound the effect of the policy. First, the COVID-19 pandemic is a potential confounder. However, the estimated responses involve more than output contraction. Regulated plants change their fuel composition, update their emission-control equipment, and use it more intensively. Furthermore, although the COVID-19 effect is known to be temporary, emissions reduction and abatement efforts persist.

Second, other environmental regulations are additional potential confounders—see Arkolakis et al. (2025) for a comprehensive review of environmental policies in South Korea. Among others, I focus on one potential confounder, which also uses size-based eligibility and covers a significant portion of targeted emissions. I consider the greenhouse gas emissions trading system and its preceding target management program. The target management system, introduced in 2010, served as a pilot for the emissions trading system launched in 2015 with quotas. Many large emitters subsequently moved to the trading system in 2015. The system has three phases: 2015–2017 (1st phase); 2018–2020 (2nd phase); and 2021–2025 (3rd phase). Because these policies differ in timing, target greenhouse gases rather than local air pollutants, and cover the whole nation, they are less likely to confound the effect of the local air pollutant cap-and-trade expansion. The effects of the policy at both district–industry and plant levels show up in 2020 and are spatial.

[32] Maintaining a balanced panel is particularly restrictive for power plants because public-private contracting and operator changes are common.

However, they also target large polluters, and a significant portion of the regulated plants overlap. Out of 940 regulated manufacturing plants, 531 plants are in the trading system and 621 are in the pilot program. Out of 153 regulated power plants, 109 are in the trading system, and 115 are in the pilot program. So I conduct several sub-sample analyses to assess whether these plants respond differently in the plant-level analysis. As shown in Tables A.16 and A.17, I find similar results. Appendix D.3 discusses the remaining policies.

D.1 District–Industry Level

First, I check if district–industry cells in the Seoul Capital Area show different patterns because this area has been covered by cap-and-trade since 2007. The effects of the earlier phase remain unsettled. Im (2023) finds reductions in ambient NO_x concentrations, whereas Lee (2019) argues that excessive permit allocations provided little incentive to reduce emissions. This paper is, to my knowledge, the first attempt to evaluate the 2020 expansion. As an experiment, I exclude observations in the Seoul Capital Area from the sample to address potential anticipatory and long-term effects and find that the results remain qualitatively similar (Tables A.7 and A.8).

Second, I exclude cells of Ongjin district, which is only partly regulated, to assess sensitivity to this partial coverage and find similar results (Table A.11). Third, I re-estimate the baseline pooled specifications with standard errors clustered at the district level to allow for correlation across sectors within districts (Tables A.12–A.13). The point estimates are unchanged. Manufacturing emissions reductions remain statistically significant, but the statistical evidence for geographic pollution shifts in manufacturing weakens. Fourth, Tables A.14–A.15 examine sensitivity to panel balance. Restricting the sample to cells with positive emissions of all three pollutants in every year yields smaller emissions declines and geographic pollution shifts in manufacturing, consistent with each other.

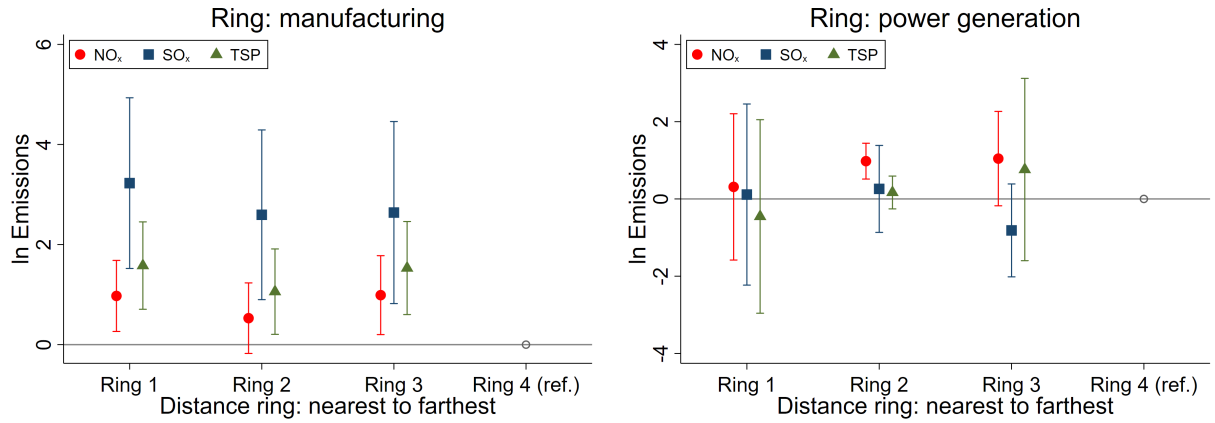
Fifth, I employ a categorical specification for the geographic pollution shift analysis, replacing the continuous measure Exposure_d with separate ring indicators. Distance is measured in kilometers between district centroids and the nearest regulated district. The sample excludes districts at or within 45 km of a regulated district. The remaining districts are grouped into distance rings defined by $\lfloor (-\ln D_d)\,/0.5 \rfloor$, where D_d denotes the nearest-regulated-district distance. The first five intervals are approximately (45,54.6], (54.6,90.0], (90.0,148.4], (148.4,244.7], and (244.7,403.4] km; more distant intervals follow the same logarithmic spacing. Let $R_{d,r}$ equal one if district d belongs to ring r and zero otherwise. Rings are numbered from nearest to farthest from regulated districts. For each sector, the reference group r_0 is the farthest ring represented in its sample. The indicator $Post_t$ equals one in 2020–2023 and zero in 2016–2019.

For manufacturing, I estimate:

$$y^p_{d,s,t} = \alpha_{d,s} + \alpha_{s,t} + \sum_{r \neq r_0} \theta_r R_{d,r} \times Post_t + \varepsilon_{d,s,t}, \tag{A.1}$$

where $y^p_{d,s,t}$ is ln Emissions, $\alpha_{d,s}$ denotes district–industry fixed effects, and $\alpha_{s,t}$ denotes industry–year fixed effects. For power generation, I replace industry–year fixed effects

Figure A.1: Geographic pollution shifts, distance rings



Notes: The left and right panels report estimates for manufacturing and power generation, respectively. Coefficients measure differences in changes in ln Emissions relative to the farthest ring. Rings are ordered from nearest to farthest. The hollow marker denotes the reference ring, normalized to zero and displayed without a confidence interval; it does not imply zero shift in that ring. The results are calculated from Equation (A.1). The 95% confidence intervals are shown jointly with the point estimates.
Sources: Ministry of Environment.

with year fixed effects, due to its small sample size. Standard errors are clustered at the district–industry level. Each coefficient θ_r measures the change in ln Emissions between the pre- and post-policy periods in ring r, relative to the corresponding change in the farthest ring. Unlike the main specification, this specification allows responses to vary freely across rings rather than imposing a linear relationship with Exposure_d.

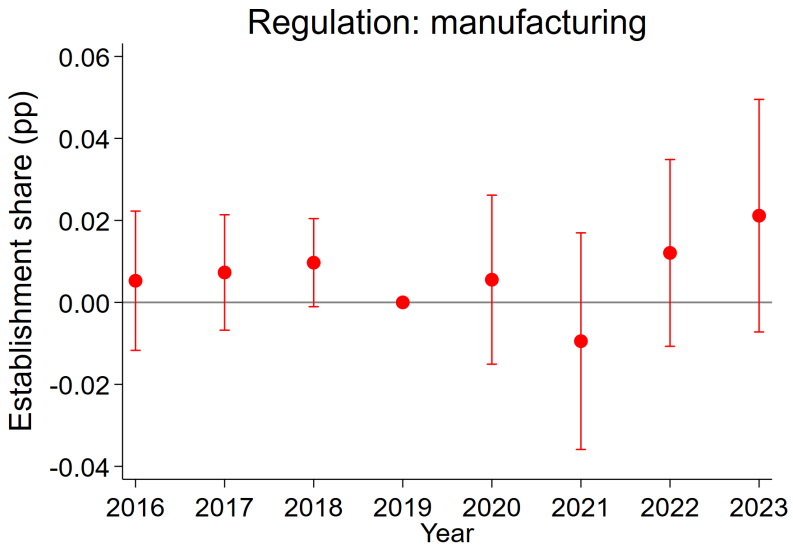
Figure A.1 presents the ring-specific estimates that are consistent with the geographic pollution shift analysis. In manufacturing, the ring-specific point estimates are positive relative to the farthest ring, and power generation estimates are mostly imprecise and do not show a consistent distance gradient.

Lastly, I examine changes in districts' shares of national manufacturing plants within three-digit industries. The outcome is the number of plants in a district–industry cell divided by the national plant count in that industry and year, expressed as a percentage. Figure A.2 shows the corresponding event study, indicating no statistically significant change in regulated districts' shares relative to unregulated districts. This supplementary analysis covers manufacturing; the much smaller fossil-power sample limits the scope for a similarly disaggregated analysis.

D.2 Plant Level

First, as above, I check if plants in the Seoul Capital Area show different patterns and find robust results (Table A.18). Second, I relax exact matching by policy area, allowing plants in other policy areas to serve as controls. This expands the pool of potential controls to include larger unregulated polluters outside the treated plants' policy areas. The results

Figure A.2: Manufacturing plant shares



Notes: The panel shows the event dynamics for manufacturing plant shares. The result is calculated from Equation (1) with three-digit KSIC industries. The 95% confidence intervals are shown jointly with the point estimates. pp denotes percentage points.
Sources: Statistics Korea.

remain similar (Table A.19).

Third, I assess sensitivity to the matching method by replacing Mahalanobis-distance matching with propensity-score matching and obtain similar estimates (Table A.23). Fourth, I investigate alternative trimming rules. The main estimates retain the same broad qualitative pattern, although equipment-use estimates are sensitive to the trimming rule (Tables A.24–A.25). Fifth, I apply the affiliated–unaffiliated comparison used for manufacturing to power plants (Table A.22). The regulated plant estimates remain imprecise, and unregulated affiliated power plants do not exhibit a statistically significant sales increase.

Sixth, Table A.21 compares unregulated manufacturing affiliates inside and outside regulated districts. Sales increase relative to unregulated unaffiliated plants in both groups. Affiliates inside regulated districts also show positive equipment-update and equipment-use responses, whereas the corresponding estimates for affiliates in unregulated districts are imprecise. These results indicate that within-firm reallocation can coexist with equipment adjustment inside regulated districts, even at plants not directly covered by the program. I interpret the later rise in within-firm leakage as evidence that abatement can initially be less costly than reallocating production within firms.

Seventh, Table A.20 examines regulated manufacturers without same-industry unregulated affiliates and unregulated affiliates in industries different from those of their regulated counterparts. These comparisons assess whether responses depend on access to same-industry affiliates and whether reallocation extends across industries. The regulated-plant responses are similar to the baseline. The sales response at different-industry affiliates is imprecisely estimated, providing no clear evidence of cross-industry sales reallocation.

Finally, I split treated plants by the number of pollutants covered by the policy to see if coverage of multiple pollutants intensifies the responses (Table A.26). Plants subject to all three pollutant-specific regulations have the largest point estimates for equipment use and electricity consumption. These results are consistent with stronger regulatory exposure inducing larger compliance responses.

D.3 Overlapping Regulations

Since 1991, the Clean Air Act in Korea has introduced several measures to curb local air pollution. The cap-and-trade program was introduced later under separate special legislation to control the amount of total emissions. First, the use of high-sulfur oil and solid fuels in certain districts has been banned. Since the policy does not vary over time, unit fixed effects address the time-invariant effect of this ban. In the baseline district–industry level specification, province–year fixed effects absorb changes common to districts within the same province. In the manufacturing plant-level analysis, I match treated and control plants within the same policy area to improve comparability in their exposure to other environmental policies.

Second, there has been a concentration regulation at the stack level and a pollution tax on the amount of emissions at the establishment level. The former establishes concentration standards for polluting sources and is enforced through command-and-control regulation. Excess emission charges are imposed on establishments violating the concentration criteria. The latter consists of basic emission charges for the establishments emitting regulated pollutants. Thus, the concentration regulation relates to the concentration, while the basic emission charge and the cap-and-trade relate to the amount of regulated emissions. These two policies have regional variation, which does not fully align with that of the cap-and-trade. They became more stringent nationwide in 2020, but those under the cap-and-trade system became exempt from the basic emission charge. Because covered plants are exempt from the basic emission charge and account for a large share of emissions in regulated districts, increases in this charge are less likely to explain the district-level emissions reductions. Also, the fixed effects absorb the common nationwide component of these changes.

These overlapping policies may make the estimated relative emissions reductions conservative. Regulated plants accounted for a large share of emissions in regulated district–industry cells, while higher pollution taxes gave unregulated plants additional incentives to reduce emissions. When I compare regulated district–industry cells to unregulated district–industry cells, I find emissions reductions. Similarly, I compare regulated plants to those that were more likely to be influenced by this overlapping policy and find abatement efforts.

Third, since 2017, there has been an attempt to integrate environmental permitting across pollution media at the establishment level and to increase its sectoral coverage over time. The integration may have tightened requirements through Best Available Techniques (BAT) standards. But the policy is nationwide, and the timing is different. Hence, the industry–year fixed effects absorb shocks common to the corresponding groups. Fourth, there have been several attempts to decrease particulate emissions since 2019. But they are

seasonal, have different timings, and their regional variations are across provinces, whose effects should be captured by the province–year fixed effects at the district–industry level.

Taken together, these considerations reduce concern that other environmental policies drive the estimated responses. However, districts and sectors may have been differently exposed to these policies. I therefore augment both district–industry specifications with interactions between mean ln Emissions over 2016–2019 and year indicators, allowing emissions trends to vary with baseline pollution levels, although these controls may also absorb policy responses that vary with baseline pollution, rather than only differential trends unrelated to the policy. I again find qualitatively similar patterns (Tables A.9–A.10).

E Theory

E.1 Full Model

This section contains the model derivations omitted from the main text.

E.1.1 Manufacturing Firms

The cost minimization problem of a plant in district l and manufacturing sector j with firm productivity ϕ, headquarters in district i, and plant productivity z can be written as:

$$\mu_{il,j}(\phi,z) = \min_{k,e,\ell;\{e^f\}_{f\in\mathbb{F}},c\in[0,1]} \left\{ rk + \sum_{f\in\mathbb{F}} p^f e^f + w_l \ell + p^e \delta c \bar{q} + \sum_{p\in\Lambda_{il,j}(\phi,z)} \tau_l^p m_{il,j}(\phi,z)\, s_{il,j}^p \right\}$$
$$\text{s.t. } \bar{q} = \phi z k^{\alpha_j^K} e^{\alpha_j^E} \ell^{\alpha_j^L}, \tag{A.2}$$
$$e = \left(\sum_{f\in\mathbb{F}} \alpha_j^f \left(e^f\right)^{\frac{\lambda_j-1}{\lambda_j}} \right)^{\frac{\lambda_j}{\lambda_j-1}},$$
$$s_{il,j}^p = (1+c)^{-\alpha^p} \left(\kappa_j^p \bar{q} + \sum_{f\in\mathbb{F}} \kappa^{p,f} e^f \right),$$

where $\bar{q} = 1$. The outer exponents satisfy $\alpha_j^K + \alpha_j^E + \alpha_j^L = 1$, with each exponent positive, and the inner weights satisfy $\sum_{f\in\mathbb{F}} \alpha_j^f = 1$ with $\lambda_j > 1$.

The first-order conditions with respect to k, e, and ℓ are:

$$\begin{aligned} r &= \frac{\alpha_j^K}{k} H_{il,j}(\phi, z), \\ p_{il,j}^E(\phi, z) &= \frac{\alpha_j^E}{e} H_{il,j}(\phi, z), \\ w_l &= \frac{\alpha_j^L}{\ell} H_{il,j}(\phi, z), \end{aligned}$$

where

$$H_{il,j}(\phi, z) = \bar{q}\left[\mu_{il,j}(\phi, z) - p^e \delta c - \sum_{p \in \Lambda_{il,j}(\phi,z)} \tau_l^p m_{il,j}(\phi, z)(1+c)^{-\alpha^p} \kappa_j^p\right].$$

Taking ratios between the factors of production, I can express all production factors in terms of labor, factor prices, and technology parameters:

$$\begin{aligned} \frac{k}{\ell} &= \frac{\alpha_j^K}{\alpha_j^L} \frac{w_l}{r}, \\ \frac{e}{\ell} &= \frac{\alpha_j^E}{\alpha_j^L} \frac{w_l}{p_{il,j}^E(\phi, z)}. \end{aligned} \tag{A.3}$$

The first-order conditions with respect to e^f and an interior solution c are:

$$\tilde{p}_{il,j}^f(\phi, z) = p_{il,j}^E(\phi, z)\, \alpha_j^f \left(\frac{e}{e^f}\right)^{\frac{1}{\lambda_j}}$$

$$p^e \delta \bar{q} = \sum_{p \in \Lambda_{il,j}(\phi,z)} \alpha^p \tau_l^p m_{il,j}(\phi, z)\left(1 + c_{il,j}(\phi, z)\right)^{-\alpha^p - 1} \left(\kappa_j^p \bar{q} + \sum_{f \in \mathbb{F}} \kappa^{p,f} e_{il,j}^f(\phi, z)\right),$$

where $\tilde{p}_{il,j}^e(\phi, z) = p^e$ and $\tilde{p}_{il,j}^f(\phi, z) = p^f + \sum_{p \in \Lambda_{il,j}(\phi,z)} \tau_l^p m_{il,j}(\phi, z)\left(1 + c_{il,j}(\phi, z)\right)^{-\alpha^p} \kappa^{p,f}$ for $f \neq e$. Dividing this condition by output, $\bar{q}$, gives:

$$p^e \delta = \sum_{p \in \Lambda_{il,j}(\phi,z)} \alpha^p \tau_l^p m_{il,j}(\phi, z)\left(1 + c_{il,j}(\phi, z)\right)^{-\alpha^p - 1} \left(\kappa_j^p + \sum_{f \in \mathbb{F}} \kappa^{p,f} \frac{e_{il,j}^f(\phi, z)}{\bar{q}}\right).$$

Thus, the benefit of using equipment intensively depends on pre-abatement emissions per unit of output, weighted by permit prices and equipment effectiveness. Under CRS, conditional fuel demands per unit of output are independent of plant scale. The equipment electricity requirement is also proportional to output, preserving a constant minimized unit cost.

Using the first-order conditions with respect to e and e^f, demand for fuel f can be derived as:

$$\frac{e^f_{il,j}(\phi,z)}{e_{il,j}(\phi,z)} = \left(\frac{\tilde{p}^f_{il,j}(\phi,z)}{p^E_{il,j}(\phi,z)\,\alpha^f_j}\right)^{-\lambda_j}. \tag{A.4}$$

With this expression, $p^E_{il,j}(\phi,z)$ can be derived from the energy aggregate:

$$p^E_{il,j}(\phi,z) = \left(\sum_{f\in\mathbb{F}} \alpha^f_j \left(\frac{\tilde{p}^f_{il,j}(\phi,z)}{\alpha^f_j}\right)^{1-\lambda_j}\right)^{\frac{1}{1-\lambda_j}}, \tag{A.5}$$

where $p^E_{il,j}(\phi,z)\,e_{il,j} = \sum_{f\in\mathbb{F}} \tilde{p}^f_{il,j}(\phi,z)\,e^f_{il,j}$.

By taking a ratio of the two first-order conditions with respect to e^f and e^e, the relative fuel demand can be written as:

$$\frac{e^f_{il,j}(\phi,z)}{e^e_{il,j}(\phi,z)} = \left(\frac{\tilde{p}^f_{il,j}(\phi,z)}{\tilde{p}^e_{il,j}(\phi,z)}\right)^{-\lambda_j}\left(\frac{\alpha^f_j}{\alpha^e_j}\right)^{\lambda_j}. \tag{A.6}$$

Then, the energy aggregate can be re-expressed only with electricity as:

$$e_{il,j}(\phi,z) = e^e_{il,j}(\phi,z)\left(\sum_{f\in\mathbb{F}} \alpha^f_j \left(\frac{\tilde{p}^f_{il,j}(\phi,z)}{\tilde{p}^e_{il,j}(\phi,z)}\right)^{1-\lambda_j}\left(\frac{\alpha^f_j}{\alpha^e_j}\right)^{\lambda_j-1}\right)^{\frac{\lambda_j}{\lambda_j-1}}.$$

The unit production cost can be derived from the first-order condition with respect to ℓ and the factor ratios (Equation (A.3)):

$$\mu_{il,j}(\phi,z) = \frac{G_{il,j}(\phi,z)}{\phi z} + p^e \delta c_{il,j}(\phi,z) + \sum_{p\in\Lambda_{il,j}(\phi,z)} \tau^p_l\, m_{il,j}(\phi,z)\left(1+c_{il,j}(\phi,z)\right)^{-\alpha^p}\kappa^p_j, \tag{A.7}$$

where $G_{il,j}(\phi,z) = \left(r/\alpha^K_j\right)^{\alpha^K_j}\left(p^E_{il,j}(\phi,z)/\alpha^E_j\right)^{\alpha^E_j}\left(w_l/\alpha^L_j\right)^{\alpha^L_j}$ is the Cobb–Douglas factor cost index. The unit cost does not depend on $q_{il,j}(\phi,z)$. Because the manufacturing technology exhibits CRS, the total operation cost equals $q_{il,j}(\phi,z)\,\mu_{il,j}(\phi,z)$.

Using the fuel demand (Equation (A.4)), the energy price (Equation (A.5)), factor ratios (Equation (A.3)), and the unit production cost (Equation (A.7)), the use of emission-control equipment can also be characterized. Let $h_{il,j}(c)$ denote the marginal benefit of equipment use minus its marginal electricity cost, both divided by the energy composite. All conditional input demands and price indices on the right-hand side are evaluated at

the candidate equipment intensity c:

$$h_{il,j}(c) = \sum_{p \in \Lambda_{il,j}(\phi,z)} \alpha^p \tau_l^p m_{il,j}(\phi,z)(1+c)^{-\alpha^p - 1} \times \left(\kappa_j^p \frac{\phi z p_{il,j}^E(\phi,z)}{\alpha_j^E G_{il,j}(\phi,z)} + \sum_f \kappa^{p,f} \left(\frac{\tilde{p}_{il,j}^f(\phi,z)}{p_{il,j}^E(\phi,z)\alpha_j^f} \right)^{-\lambda_j} \right) - p^e \delta \frac{\phi z p_{il,j}^E(\phi,z)}{\alpha_j^E G_{il,j}(\phi,z)}. \tag{A.8}$$

The Cobb–Douglas input conditions imply $q_{il,j}/e_{il,j} = \phi z p_{il,j}^E / \left(\alpha_j^E G_{il,j}\right)$. Because $h_{il,j}(c)$ need not be monotone, the global optimum is obtained by comparing the minimized unit cost at $c = 0$, $c = 1$, and every interior $c \in (0,1)$ satisfying $h_{il,j}(c) = 0$. The optimal $c_{il,j}^*$ is the candidate that yields the lowest unit cost. For an unregulated plant, $c_{il,j}^* = 0$.

Then, the plant-level variety price combines a markup, a multiregional production friction, a unit production cost, and an iceberg trade cost:

$$p_{iln,j}(\phi,z) = \mathcal{M}_j \gamma_{il,j} \mu_{il,j}(\phi,z)\, t_{ln,j}.$$

From the household demand structure (Equation (11)), the revenue of this plant can be derived as:

$$\begin{aligned} R_{il,j}(\phi,z) &= \sum_n p_{iln,j}(\phi,z)\, q_{iln,j}(\phi,z) \\ &= \sum_n \left(\frac{p_{iln,j}(\phi,z)}{P_{in,j}(\phi)} \right)^{1-\eta_j} \left(\frac{P_{in,j}(\phi)}{P_{n,j}} \right)^{1-\sigma_j} X_{n,j}, \end{aligned} \tag{A.9}$$

and the quantity demand for this plant can be derived as:

$$q_{il,j}(\phi,z) = \gamma_{il,j} \sum_n t_{ln,j} \left(\frac{p_{iln,j}(\phi,z)}{P_{in,j}(\phi)} \right)^{-\eta_j} \left(\frac{P_{in,j}(\phi)}{P_{n,j}} \right)^{-\sigma_j} \frac{X_{n,j}}{P_{n,j}}, \tag{A.10}$$

where $q_{il,j}(\phi,z) = \gamma_{il,j} \sum_n t_{ln,j} q_{iln,j}(\phi,z)$. The corresponding profit is:

$$\pi_{il,j}(\phi,z) = \frac{R_{il,j}(\phi,z)}{\sigma_j} + \sum_{p \in \Lambda_{il,j}(\phi,z)} \tau_l^p A_{il,j}^p(\phi,z), \tag{A.11}$$

where $A_{il,j}^p(\phi,z)$ denotes the initial allocation of permits for pollutant p to the manufacturing plant. It is determined from pre-regulation emissions under the policy-specific allocation rule and is independent of contemporaneous production and abatement choices.

Pricing and Markups I derive the pricing rule under centralized firm-level pricing and then contrast it with decentralized plant-level pricing. Throughout, input prices, permit prices, and the set of plant varieties are given. Each firm is atomistic in the sectoral market and takes $P_{n,j}$ and $Q_{n,j}$ as given. I suppress productivity and variety arguments except where needed. Let $\mathrm{MC}_{iln,j} \equiv \gamma_{il,j}\mu_{il,j}t_{ln,j}$ denote the marginal cost of delivering a plant variety to destination n. Since plant-level unit costs are independent of output, the minimum cost of supplying $Q_{in,j}$ units of the firm composite is $\mu^F_{in,j}Q_{in,j}$:

$$\mu^F_{in,j} = \left(\sum_l \int_{\tilde{\omega} \in \Omega_{iln,j}(\omega)} \left[\mathrm{MC}_{iln,j}\left(\tilde{\omega}\right) \right]^{1-\eta_j} d\tilde{\omega} \right)^{\frac{1}{1-\eta_j}}.$$

Under centralized pricing, the firm internalizes substitution among its plant varieties. Its destination-specific pricing problem can therefore be written as:

$$\max_{P_{in,j}>0} \left(P_{in,j} - \mu^F_{in,j} \right) Q_{in,j}, \qquad Q_{in,j} = \left(\frac{P_{in,j}}{P_{n,j}} \right)^{-\sigma_j} Q_{n,j}.$$

The first-order condition gives:

$$0 = Q_{in,j} - \sigma_j \left(P_{in,j} - \mu^F_{in,j} \right) \frac{Q_{in,j}}{P_{in,j}} \quad \Longrightarrow \quad P_{in,j} = \frac{\sigma_j}{\sigma_j - 1} \mu^F_{in,j}.$$

The cost-minimizing allocation across plant varieties and household demand must coincide:

$$q_{iln,j} = \left(\frac{\mathrm{MC}_{iln,j}}{\mu^F_{in,j}} \right)^{-\eta_j} Q_{in,j} = \left(\frac{p_{iln,j}}{P_{in,j}} \right)^{-\eta_j} Q_{in,j}.$$

Consequently, $p_{iln,j}/P_{in,j} = \mathrm{MC}_{iln,j}/\mu^F_{in,j}$, implying:

$$p_{iln,j} = \frac{\sigma_j}{\sigma_j - 1} \mathrm{MC}_{iln,j}.$$

For comparison, under decentralized pricing, an atomistic plant takes its firm's price index $P_{in,j}$ and composite demand $Q_{in,j}$ as given. Its demand elasticity is then η_j. Maximizing $\left(p_{iln,j} - \mathrm{MC}_{iln,j} \right) q_{iln,j}$ gives:

$$0 = q_{iln,j} - \eta_j \left(p_{iln,j} - \mathrm{MC}_{iln,j} \right) \frac{q_{iln,j}}{p_{iln,j}} \quad \Longrightarrow \quad p_{iln,j} = \frac{\eta_j}{\eta_j - 1} \mathrm{MC}_{iln,j}.$$

Thus, the markup depends on which substitution margins the pricing decision internalizes. Under the centralized pricing assumed in this paper, σ_j determines the common markup, while η_j determines the allocation across plant varieties and the corresponding CES unit cost.

E.1.2 Power Plants

Conditional on dispatched output y, the power plant's cost-minimization problem is:

$$V_{i,g}(y,z) = \min_{e_g;a\in[0,1)} \left\{ p^g e_g + \sum_{p\in\Lambda_{i,g}(z)} \tau_i^p m_{i,g}(z)\, s_{i,g}^p(y,z) \right\}$$
$$\text{s.t. } y = (1-a)\min\left\{k_{i,g}(z), z e_g\right\},$$
$$s_{i,g}^p(y,z) = (1-a)^{\frac{1}{\xi^p}} \kappa_g^p e_g.$$

When capacity is slack, so that $z e_g < k_{i,g}(z)$, the first-order conditions with respect to an interior solution $(1-a)$ and e_g are:

$$\sum_{p\in\Lambda_{i,g}(z)} \tau_i^p m_{i,g}(z) \frac{1}{\xi^p} (1-a)^{\frac{1}{\xi^p}-1} \kappa_g^p e_g = \nu_{i,g}(y,z)\, z e_g,$$

$$p^g + \sum_{p\in\Lambda_{i,g}(z)} \tau_i^p m_{i,g}(z) (1-a)^{\frac{1}{\xi^p}} \kappa_g^p = \nu_{i,g}(y,z)\, z\,(1-a),$$

where $\nu_{i,g}(y,z)$ is the marginal generation cost.

Define the slack-capacity first-order-condition residual as:

$$H_{i,g}(1-a) = \sum_{p\in\Lambda_{i,g}(z)} \tau_i^p m_{i,g}(z) \left(\frac{1}{\xi^p} - 1\right) (1-a)^{\frac{1}{\xi^p}} \kappa_g^p - p^g. \tag{A.12}$$

Because $H_{i,g}(1-a)$ is increasing in $1-a$, the slack-capacity choice is $a_{i,g}^0(z) = 0$ when $H_{i,g}(1) \leq 0$. If $p^g > 0$ and $H_{i,g}(1) > 0$, there is a unique interior solution satisfying $H_{i,g}\left(1 - a_{i,g}^0(z)\right) = 0$. If $p^g = 0$ and $\tau_i^p \kappa_g^p > 0$ for at least one $p \in \Lambda_{i,g}(z)$, the slack-capacity cost expression, considered without the capacity constraint, has no finite minimizer: with free fuel, the pollution cost approaches zero as $1-a$ approaches zero. I therefore define the limiting slack-capacity choice by $1 - a_{i,g}^0(z) = 0$. Imposing the capacity constraint gives the cost-minimizing equipment choice:

$$1 - a_{i,g}(y,z) = \max\left\{1 - a_{i,g}^0(z), \frac{y}{k_{i,g}(z)}\right\}.$$

When $1 - a_{i,g}^0(z) = 0$, every positive dispatched output lies in the capacity-bound branch. If $p^g = 0$ but $\tau_i^p \kappa_g^p = 0$ for every p, I set $a_{i,g}^0(z) = 0$. An unregulated plant likewise has $a_{i,g}^0(z) = 0$.

Fuel g use is:

$$e_{i,g}(y,z) = \frac{y}{z\left(1 - a_{i,g}(y,z)\right)}.$$

When capacity binds, $1 - a_{i,g}(y,z) = y/k_{i,g}(z)$, and hence $e_{i,g}(y,z) = k_{i,g}(z)/z$. The total generation cost is:

$$V_{i,g}(y,z) = \frac{p^g k_{i,g}(z)}{z} + \frac{k_{i,g}(z)}{z} \sum_{p \in \Lambda_{i,g}(z)} \tau_i^p m_{i,g}(z) \kappa_g^p \left(\frac{y}{k_{i,g}(z)} \right)^{1/\xi^p}.$$

The fuel-cost term is fixed with respect to y. Thus, p^g enters the total cost but drops out of the capacity-bound branch of the marginal cost, which reflects the marginal pollution cost of relaxing emission control.

Then, the marginal generation cost is:

$$\nu_{i,g}(y,z) = \frac{1}{\left(1 - a_{i,g}^0(z)\right)} \frac{\left(p^g + \sum_{p \in \Lambda_{i,g}(z)} \tau_i^p m_{i,g}(z) \left(1 - a_{i,g}^0(z)\right)^{\frac{1}{\xi^p}} \kappa_g^p \right)}{z},$$

if $0 < y \leq \left(1 - a_{i,g}^0(z)\right) k_{i,g}(z)$ and

$$\nu_{i,g}(y,z) = \frac{1}{z} \sum_{p \in \Lambda_{i,g}(z)} \frac{\tau_i^p m_{i,g}(z) \kappa_g^p}{\xi^p} \left(\frac{y}{k_{i,g}(z)} \right)^{\frac{1}{\xi^p} - 1}, \tag{A.13}$$

if $\left(1 - a_{i,g}^0(z)\right) k_{i,g}(z) < y \leq k_{i,g}(z)$. At zero output, I set $V_{i,g}(0,z) = e_{i,g}(0,z) = s_{i,g}^p(0,z) = 0$ and define $\nu_{i,g}(0,z)$ by its right-hand limit. In the capacity-bound branch, $\nu_{i,g}(y,z) = \partial V_{i,g}(y,z)/\partial y$ differs from average unit cost $V_{i,g}(y,z)/y$. This setting contrasts with manufacturing, where CRS without a capacity constraint implies that unit and marginal cost coincide.

The corresponding profit is:

$$\pi_{i,g}(z) = p^e y_{i,g}(z) - \int_0^{y_{i,g}(z)} \nu_{i,g}(u,z)\, du + \sum_{p \in \Lambda_{i,g}(z)} \tau_i^p A_{i,g}^p(z), \tag{A.14}$$

where $A_{i,g}^p(z)$ denotes the initial allocation of permits for pollutant p to the power plant. It is determined before power generation and does not affect generation or abatement choices.

E.1.3 Aggregation and Equilibrium Conditions

The price index of manufacturing sector j in district n satisfies the following condition:

$$P_{n,j} = \left(\sum_i M_{i,j} \mathbb{E}_\phi \left[\left(P_{in,j}^{1-\sigma_j}(\phi) \right) \right] \right)^{\frac{1}{1-\sigma_j}},$$

where

$$P_{in,j}(\phi) = \left(\sum_l \mathbb{E}_z \left[p_{iln,j}^{1-\eta_j}(\phi, z) \right] \right)^{\frac{1}{1-\eta_j}}.$$

The labor market clears in each district:

$$L_l = \sum_{i,j} M_{i,j} \mathbb{E}_{\phi,z} \left[\ell_{il,j}(\phi, z) \right].$$

Using the demand function (Equation (A.10)), the production function (Equation (A.2)), and the ratios of the first-order conditions (Equation (A.3)), the labor demand can be expressed as:

$$\ell_{il,j}(\phi, z) = \frac{\alpha_j^L G_{il,j}(\phi, z)}{\phi z w_l} q_{il,j}(\phi, z).$$

Final expenditure in district l and sector j is:

$$X_{l,j} = \gamma_j^C (w_l L_l + r K_l + \Pi_l).$$

Total electricity demand in manufacturing equals the sum of industrial electricity demand:

$$ED = \sum_{i,l} \sum_j ED_{il,j}, \tag{A.15}$$

where industrial electricity demand is:

$$ED_{il,j} = M_{i,j} \mathbb{E}_{\phi,z} \left[e_{il,j}^e(\phi, z) + \delta c_{il,j}(\phi, z) q_{il,j}(\phi, z) \right].$$

Using the fuel demand function (Equation (A.4)), I obtain:

$$e_{il,j}^e(\phi, z) = \left(\frac{\tilde{p}_{il,j}^e(\phi, z)}{p_{il,j}^E(\phi, z) \alpha_j^e} \right)^{-\lambda_j} \frac{\alpha_j^E G_{il,j}(\phi, z)}{\phi z p_{il,j}^E(\phi, z)} q_{il,j}(\phi, z).$$

The emissions of pollutant p from manufacturing plants in district l and sector j, $S_{l,j}^p$, can be expressed as:

$$S_{l,j}^p = \sum_i M_{i,j} \mathbb{E}_{\phi,z} \left[\left(1 + c_{il,j}(\phi, z)\right)^{-\alpha^p} \left(\kappa_j^p q_{il,j}(\phi, z) + \sum_{f \in \mathbb{F}} \kappa^{p,f} e_{il,j}^f(\phi, z) \right) \right].$$

Also, the emissions of pollutant p from power plants of fuel type g in district l, $S_{l,g}^p$, can be written as:

$$S_{l,g}^p = M_{l,g} \mathbb{E}_z \left[s_{l,g}^p \left(y_{l,g}(z; p^e), z \right) \right],$$

where $s_{l,g}^p(0, z) = 0$. Then, the total emissions are the sum of these sectoral emissions:

$S_l^p = \sum_j S_{l,j}^p + \sum_g S_{l,g}^p$.

I abstract from intertemporal banking because it is limited to one year, and from borrowing because it is prohibited. Permit-market clearing requires regulated emissions not to exceed the cap. Permit prices are non-negative, and a positive permit price requires the cap to bind. In equilibrium, unused permits require a zero permit price. Thus, the static permit-market clearing conditions for the cap-and-trade model are as follows:

$$0 \leq \tau_a^p \;\perp\; C_a^p - \sum_{l \in a} \left(\sum_j \bar{S}_{l,j}^p + \sum_g \bar{S}_{l,g}^p \right) \geq 0,$$

where

$$\bar{S}_{l,j}^p = \sum_i M_{i,j} \mathbb{E}_{\phi,z} \left[\mathbf{1} \left\{ p \in \Lambda_{il,j}(\phi, z) \right\} s_{il,j}^p(\phi, z) \right],$$

and

$$\bar{S}_{l,g}^p = M_{l,g} \mathbb{E}_z \left[\mathbf{1} \left\{ p \in \Lambda_{l,g}(z) \right\} s_{l,g}^p \left(y_{l,g}(z; p^e), z \right) \right].$$

Initial permit allocations separately exhaust the cap:

$$C_a^p = \sum_{l \in a} \left(\sum_j \bar{A}_{l,j}^p + \sum_g \bar{A}_{l,g}^p \right),$$

where $\bar{A}_{l,j}^p$ and $\bar{A}_{l,g}^p$ are the corresponding aggregates of plant-level initial allocations. Subsequent permit trades only reallocate these permits across regulated plants.

E.2 Proposition Proofs

Symbols defined in this subsection are local to it and do not carry their main-text meanings. I suppress firm, plant, sector, and productivity indices, absorb each fixed plant-specific compliance-cost multiplier into its effective pollution prices, and evaluate all derivatives at a regular interior manufacturing equilibrium, with $c_i \in (0,1)$ and $A_i \equiv 1 + c_i$. Let Λ_i denote the set of active pollutants at the plant. For each $r \in \Lambda_i$, let $B_i^r \equiv \kappa^r q_i + \sum_{f \in \mathbb{F}} \kappa^{r,f} e_i^f$ denote emissions before equipment use and $\mathcal{B}_i^r \equiv B_i^r / q_i$ the corresponding emissions per unit of output. Partial derivatives with respect to τ_i^p hold A_i and all other permit prices fixed. For $f \neq e$, define the inner-nest fuel and electricity expenditure shares by $s_i^f \equiv \tilde{p}_i^f e_i^f / \left(p_i^E e_i \right)$ and $s_i^e \equiv \tilde{p}_i^e e_i^e / \left(p_i^E e_i \right) = 1 - \sum_{f \neq e} s_i^f$. The process and fuel-specific shares of these emissions are $\Omega_i^{r,q} \equiv \kappa^r q_i / B_i^r$ and $\Omega_i^{r,f} \equiv \kappa^{r,f} e_i^f / B_i^r$. The share of fuel f's effective price attributable to pollutant r is $\chi_i^{r,f} \equiv \tau_i^r A_i^{-\alpha^r} \kappa^{r,f} / \tilde{p}_i^f$. Define $S_i^r \equiv \sum_{f \neq e} s_i^f \chi_i^{r,f}$, $T_i^f \equiv \sum_{r \in \Lambda_i} \alpha^r \chi_i^{r,f}$, and $S_i^\alpha \equiv \sum_{f \neq e} s_i^f T_i^f$.

The conditions below are evaluated at the reference equilibrium using the expressions derived for equipment use, effective fuel prices, and demand responses. They characterize the signs of local comparative statics.

Proof of Proposition 1. The outer Cobb–Douglas first-order conditions give $e_i/q_i = \alpha^E G_i/\left(\phi z p_i^E\right)$ and $d\ln G_i = \alpha^E d\ln p_i^E$ at fixed non-energy input prices. The inner CES condition gives $e_i^e/e_i = \left(p^e/\left(p_i^E\alpha^e\right)\right)^{-\lambda}$. Consequently, $d\ln\left(e_i^e/q_i\right) = \left(\lambda - 1 + \alpha^E\right) d\ln p_i^E$. The Cobb–Douglas and CES first-order conditions imply:

$$\begin{aligned}
\left.\frac{\partial \ln\left(e_i^f/e_i^e\right)}{\partial \ln \tau_i^p}\right|_{A_i} &= -\lambda\chi_i^{p,f},\\
\left.\frac{\partial \ln\left(q_i/e_i^e\right)}{\partial \ln \tau_i^p}\right|_{A_i} &= -\left(\lambda - 1 + \alpha^E\right) S_i^p,\\
\left.\frac{\partial \ln\left(e_i^f/e_i^e\right)}{\partial \ln A_i}\right|_{\tau_i^p} &= \lambda T_i^f,\\
\left.\frac{\partial \ln\left(q_i/e_i^e\right)}{\partial \ln A_i}\right|_{\tau_i^p} &= \left(\lambda - 1 + \alpha^E\right) S_i^\alpha.
\end{aligned}$$

Subtracting the output-to-electricity response from the fuel-to-electricity response gives $\left.\frac{\partial \ln\left(e_i^f/q_i\right)}{\partial \ln \tau_i^p}\right|_{A_i} = -\lambda\chi_i^{p,f} + \left(\lambda - 1 + \alpha^E\right) S_i^p$ and $\left.\frac{\partial \ln\left(e_i^f/q_i\right)}{\partial \ln A_i}\right|_{\tau_i^p} = \lambda T_i^f - \left(\lambda - 1 + \alpha^E\right) S_i^\alpha$. Since $\mathcal{B}_i^r = \kappa^r + \sum_{f\neq e}\kappa^{r,f} e_i^f/q_i$, its log derivative is the corresponding emissions-share-weighted sum of these fuel-intensity derivatives. Define:

$$K_i^{r,p} \equiv -\left.\frac{\partial \ln \mathcal{B}_i^r}{\partial \ln \tau_i^p}\right|_{A_i}, \qquad L_i^r \equiv \left.\frac{\partial \ln \mathcal{B}_i^r}{\partial \ln A_i}\right|_{\tau_i^p}.$$

Then,

$$K_i^{r,p} = \lambda \sum_{f\neq e} \Omega_i^{r,f}\chi_i^{p,f} - \left(1 - \Omega_i^{r,q}\right)\left(\lambda - 1 + \alpha^E\right) S_i^p,$$

and

$$L_i^r = \lambda \sum_{f\neq e} \Omega_i^{r,f} T_i^f - \left(1 - \Omega_i^{r,q}\right)\left(\lambda - 1 + \alpha^E\right) S_i^\alpha = \sum_{r'\in\Lambda_i} \alpha^{r'} K_i^{r,r'}.$$

The first-order condition with respect to the common emission-control equipment is $p^e\delta q_i = \sum_{r\in\Lambda_i}\alpha^r\tau_i^r A_i^{-\alpha^r-1}B_i^r$. Equivalently, $p^e\delta = A_i^{-1}\sum_{r\in\Lambda_i}\alpha^r\tau_i^r A_i^{-\alpha^r}\mathcal{B}_i^r$. Let $\pi_i^r \equiv \frac{\alpha^r\tau_i^r A_i^{-\alpha^r}\mathcal{B}_i^r}{\sum_{m\in\Lambda_i}\alpha^m\tau_i^m A_i^{-\alpha^m}\mathcal{B}_i^m}$, $\bar{\alpha}_i \equiv \sum_{r\in\Lambda_i}\pi_i^r\alpha^r$, $\bar{K}_i^p \equiv \sum_{r\in\Lambda_i}\pi_i^r K_i^{r,p}$, and $\bar{L}_i \equiv \sum_{r\in\Lambda_i}\pi_i^r L_i^r$. Taking log derivatives of the first-order condition with respect to the pollution price gives:

$$\varepsilon_i^{A,p} \equiv \frac{d\ln A_i}{d\ln \tau_i^p} = \frac{\pi_i^p - \bar{K}_i^p}{1 + \bar{\alpha}_i - \bar{L}_i}.$$

To map this response into the accounting decomposition, first note that fixed production-electricity prices and the Cobb–Douglas and CES input conditions imply $d\ln\left(q_i/e_i^e\right) =$

$\left[1-\left(1-\alpha^E\right)/\lambda\right]\sum_{f\neq e}s_i^f d\ln\left(e_i^f/e_i^e\right)$. Let $\widetilde{e}_i^e \equiv e_i^e + \delta c_i q_i$ denote total electricity use. Define $\omega_i^e \equiv e_i^e/\widetilde{e}_i^e$ and $\varepsilon_i^e \equiv \left.\frac{\partial \ln \widetilde{e}_i^e}{\partial \ln A_i}\right|_{q_i, e_i^e} = \delta A_i q_i/\widetilde{e}_i^e$. Differentiating total electricity use gives $d\ln \widetilde{e}_i^e = \omega_i^e d\ln e_i^e + \left(1-\omega_i^e\right) d\ln q_i + \varepsilon_i^e d\ln A_i$. Using these identities to express production and fuel use relative to total electricity, define:

$$\mathcal{W}_i^{p,f} \equiv \Omega_i^{p,f} + \Omega_i^{p,q}\frac{\omega_i^e\left(\lambda-1+\alpha^E\right)}{\lambda-\left(1-\omega_i^e\right)\left(\lambda-1+\alpha^E\right)\left(1-s_i^e\right)}s_i^f,$$

and

$$\begin{aligned}\mathcal{C}_i^p &\equiv \sum_{f\neq e}\mathcal{W}_i^{p,f} - 1\\ &= -\Omega_i^{p,q}\frac{\lambda-\left(\lambda-1+\alpha^E\right)\left(1-s_i^e\right)}{\lambda-\left(1-\omega_i^e\right)\left(\lambda-1+\alpha^E\right)\left(1-s_i^e\right)} \leq 0.\end{aligned}$$

Since $\lambda > 1$ and $0 < \alpha^E < 1$, $0 < \lambda - 1 + \alpha^E < \lambda$. Hence,

$$\lambda - \left(1-\omega_i^e\right)\left(\lambda-1+\alpha^E\right)\left(1-s_i^e\right) \geq 1-\alpha^E > 0.$$

The denominators are therefore positive, and the weights $\mathcal{W}_i^{p,f}$ are non-negative. Then,

$$d\ln s_i^p = \left[-\alpha^p + \mathcal{C}_i^p\varepsilon_i^e\right]d\ln A_i + \sum_{f\neq e}\mathcal{W}_i^{p,f}d\ln\left(\frac{e_i^f}{\widetilde{e}_i^e}\right) + d\ln\left(\frac{\widetilde{e}_i^e}{R_i}\right) + d\ln R_i.$$

The energy-price response is $g_i^{E,p} \equiv \frac{d\ln p_i^E}{d\ln \tau_i^p} = \sum_{f\neq e}s_i^f\left(\chi_i^{p,f} - T_i^f\varepsilon_i^{A,p}\right) = S_i^p - S_i^\alpha\varepsilon_i^{A,p}$. The Cobb–Douglas and CES conditions imply

$$\frac{d\ln\left(e_i^e/q_i\right)}{d\ln\tau_i^p} = \left(\lambda-1+\alpha^E\right)g_i^{E,p},$$

and

$$\frac{d\ln\left(e_i^f/e_i^e\right)}{d\ln\tau_i^p} = -\lambda\left(\chi_i^{p,f} - T_i^f\varepsilon_i^{A,p}\right).$$

Combining these responses with the total-electricity identity gives the three within-plant elasticities:

$$\begin{aligned}\Gamma_i^p &= -\left[\alpha^p - \mathcal{C}_i^p\varepsilon_i^e\right]\varepsilon_i^{A,p},\\ \psi_i^{p,f} &= -\lambda\left(\chi_i^{p,f} - T_i^f\varepsilon_i^{A,p}\right) + \left(1-\omega_i^e\right)\left(\lambda-1+\alpha^E\right)g_i^{E,p} - \varepsilon_i^e\varepsilon_i^{A,p},\\ \Xi_i^p &= \omega_i^e\left(\lambda-1+\alpha^E\right)g_i^{E,p} + \varepsilon_i^e\varepsilon_i^{A,p} - \vartheta_i^p,\end{aligned}$$

where the envelope theorem and constant-markup pricing give:

$$\vartheta_i^p \equiv \frac{d \ln p_i}{d \ln \tau_i^p} = \tau_i^p A_i^{-\alpha^p} B_i^p / (\mu_i q_i) > 0.$$

Since $\mathcal{C}_i^p \leq 0$, $\alpha^p > 0$, and $\varepsilon_i^e > 0$, $\Gamma_i^p < 0$ whenever $\varepsilon_i^{A,p} > 0$. The expression for $\varepsilon_i^{A,p}$ is positive if $\pi_i^p > \bar{K}_i^p$ and $1 + \bar{\alpha}_i > \bar{L}_i$. The first condition requires the higher permit price to raise the marginal benefit of equipment after fuel use adjusts. The second ensures that this marginal benefit declines with equipment intensity, satisfying the strict second-order condition.

Holding equipment use fixed, the higher permit price encourages substitution away from fossil fuels. Additional equipment uses electricity but also reduces the pollution charge per unit of fuel. Accounting for both effects and the adjustment in production electricity, fossil-fuel use relative to total electricity falls, $\psi_i^{p,f} < 0$, when:

$$\lambda \chi_i^{p,f} + \varepsilon_i^e \varepsilon_i^{A,p} > \lambda T_i^f \varepsilon_i^{A,p} + (1 - \omega_i^e) \left(\lambda - 1 + \alpha^E \right) g_i^{E,p},$$

for every $f \neq e$.

Finally, the production-electricity response combines two forces. With electricity prices fixed, a higher energy price index induces substitution toward electricity, contributing $\lambda g_i^{E,p}$, but reduces total energy use per unit of output, contributing $-\left(1 - \alpha^E\right) g_i^{E,p}$. Their combined contribution to total electricity intensity is weighted by the production-electricity share, ω_i^e. Additional equipment use contributes $\varepsilon_i^e \varepsilon_i^{A,p}$ through the electricity required for abatement. By contrast, the output-price increase, ϑ_i^p, raises revenue per unit of output. Electricity use per dollar of revenue therefore rises, $\Xi_i^p > 0$, if and only if the net increase in electricity requirements per unit of output exceeds this price effect:

$$\omega_i^e \left(\lambda - 1 + \alpha^E \right) g_i^{E,p} + \varepsilon_i^e \varepsilon_i^{A,p} > \vartheta_i^p.$$

□

Proof of Proposition 2. Consider the regulated plant in district i, belonging to firm ω, and fix a destination n and sector j. By the envelope theorem and constant-markup pricing,

$$\vartheta_i^p \equiv \frac{d \ln p_i}{d \ln \tau_i^p} = \frac{\tau_i^p A_i^{-\alpha^p} B_i^p}{\mu_i q_i} > 0.$$

In this proof, i and l denote plant locations. I suppress the headquarters-location subscript and write $P_{n,j}(\omega)$ for firm ω's price index in destination n, while $P_{n,j}$ denotes the sectoral price index. Define $v_{\omega n,j}^p \equiv \frac{d \ln P_{n,j}(\omega)}{d \ln \tau_i^p}$ and $v_{n,j}^p \equiv \frac{d \ln P_{n,j}}{d \ln \tau_i^p}$. CES price-index derivatives are expenditure-share weighted averages of component price elasticities. The relevant local exposure condition is $\vartheta_i^p \geq v_{\omega n,j}^p > v_{n,j}^p > 0$. Under this ordering, the plant's price rises at

least as much as its firm's price index, so the plant does not gain expenditure share within the firm. The firm's price index rises more than the sectoral index, shifting expenditure toward competing firms. Together, these relative-price changes reduce the plant's revenue at that destination. The nested CES revenue equation implies:

$$\frac{d \ln R_{in,j}}{d \ln \tau_i^p} = -\left(\eta_j - 1\right)\left(\vartheta_i^p - v_{\omega n,j}^p\right) - \left(\sigma_j - 1\right)\left(v_{\omega n,j}^p - v_{n,j}^p\right) < 0.$$

The aggregate revenue elasticity of the regulated plant is:

$$\Upsilon_i^p \equiv \frac{d \ln R_i}{d \ln \tau_i^p} = \sum_n \frac{R_{in,j}}{R_i} \frac{d \ln R_{in,j}}{d \ln \tau_i^p}.$$

For an unregulated plant l in the same firm, the nested CES demand equation gives:

$$\frac{d \ln R_{ln,j}}{d \ln \tau_i^p} = \left(\eta_j - \sigma_j\right) v_{\omega n,j}^p + \left(\sigma_j - 1\right) v_{n,j}^p.$$

Both terms are positive because $\eta_j > \sigma_j > 1$ and $v_{\omega n,j}^p > v_{n,j}^p > 0$. The affiliate gains demand from the firm's regulated plants, and this gain exceeds the loss of demand to competing firms.

For an unregulated plant in a different firm whose own firm price index is unaffected,

$$\frac{d \ln R_{ln,j}}{d \ln \tau_i^p} = \left(\sigma_j - 1\right) v_{n,j}^p > 0,$$

at a destination with positive price-index exposure. If that firm also owns regulated plants, its price-index response enters through the additional term $\left(\eta_j - \sigma_j\right) v_{\omega' n,j}^p$, where ω' denotes that firm.

For the regulated plant, summing over destinations preserves the negative revenue response provided the stated exposure ordering holds at each destination it serves. For an unregulated plant, a sufficient local condition for positive total revenue growth is that the destination-specific response above be non-negative at every destination it serves and strictly positive at at least one destination with positive baseline revenue. The unregulated plant faces unchanged own input prices and permit prices, so its optimal equipment intensity and emissions per dollar of revenue are unchanged. Constant returns to scale then imply:

$$\frac{d \ln s_l^p}{d \ln \tau_i^p} = \frac{d \ln R_l}{d \ln \tau_i^p} > 0.$$

This proves the market-mediated leakage result under the stated local exposure conditions. □

F Model Estimation

F.1 Production Function Estimation

The production function is estimated using pre-regulation data. I drop j for simplicity because the production function is estimated for each sector separately. After the realization of ϕ and z, the revenue of a plant with (ϕ, z) headquartered in district i and producing in district l is the sum of revenues across all districts n:

$$\begin{aligned} R_{il}(\phi,z) &= \sum_n p_{iln}(\phi,z)\, q_{iln}(\phi,z) \\ &= \sum_n Q_n^{\frac{1}{\eta}} P_n^{\frac{\sigma}{\eta}} P_{in}^{\frac{\eta-\sigma}{\eta}}(\phi)\, q_{iln}^{\frac{\eta-1}{\eta}}(\phi,z) \\ &= q_{il}^{\frac{\eta-1}{\eta}}(\phi,z) \sum_n Q_n^{\frac{1}{\eta}} P_n^{\frac{\sigma}{\eta}} P_{in}^{\frac{\eta-\sigma}{\eta}}(\phi)\, s_{iln}^{\frac{\eta-1}{\eta}}(\phi) \\ &= q_{il}^{\frac{\eta-1}{\eta}}(\phi,z)\, \Gamma_{il}(\phi), \end{aligned}$$

where the total quantity manufactured by the plant is $q_{il}(\phi,z) = \gamma_{il} \sum_n t_{ln} q_{iln}(\phi,z)$, the ratio of output delivered to district n to total output produced is

$$s_{iln}(\phi,z) = \frac{q_{iln}(\phi,z)}{q_{il}(\phi,z)} = s_{iln}(\phi),$$

and $\Gamma_{il}(\phi) = \sum_n Q_n^{\frac{1}{\eta}} P_n^{\frac{\sigma}{\eta}} P_{in}^{\frac{\eta-\sigma}{\eta}}(\phi)\, s_{iln}^{\frac{\eta-1}{\eta}}(\phi)$. I use the unit production cost (Equation (7)), the demand function (Equation (11)), and the price index (Equation (10)). This result shows that the ratio is independent of z.

I now derive the estimating equation from the plant's electricity first-order condition. I suppress the plant, district, and sector indices and write $\chi = \phi z$. Let R^* denote revenue implied by the maintained demand relation at the input-choice stage. The preceding demand derivation gives:

$$R^* = \Gamma \left(\chi k^{\alpha^K} e^{\alpha^E} \ell^{\alpha^L}\right)^{\frac{\eta-1}{\eta}}, \qquad e = \left[\sum_{f \in \mathbb{F}} \alpha^f \left(e^f\right)^{\frac{\lambda-1}{\lambda}}\right]^{\frac{\lambda}{\lambda-1}}.$$

The demand shifter Γ retains the plant-specific market access and demand components derived above; it is not assumed to be common across plants. In the pre-regulation economy, the firm jointly chooses prices for its plant varieties, and each plant minimizes cost conditional on output. Centralized pricing implies that each plant's variable cost equals $\frac{\sigma-1}{\sigma} R^*$.

First, Cobb–Douglas cost minimization allocates this variable cost across factors:

$$\frac{rk}{R^*} = \frac{\sigma-1}{\sigma}\alpha^K, \qquad \frac{w\ell}{R^*} = \frac{\sigma-1}{\sigma}\alpha^L, \qquad \frac{\sum_{f \in \mathbb{F}} p^f e^f}{R^*} = \frac{\sigma-1}{\sigma}\alpha^E.$$

Since $\alpha^K + \alpha^E + \alpha^L = 1$, each exponent equals the corresponding share of $\mathcal{S} = \mathcal{S}^K + \mathcal{S}^L + \mathcal{S}^E$. With common exponents within a sector, aggregation gives:

$$\alpha^K = \frac{\mathcal{S}^K}{\mathcal{S}}, \qquad \alpha^E = \frac{\mathcal{S}^E}{\mathcal{S}}, \qquad \alpha^L = 1 - \alpha^K - \alpha^E.$$

The sums run over plant-year observations within the sector. These restrictions require cost-minimizing factor choices, including capital evaluated at its rental or user cost. A capital-stock value cannot be substituted for rk.

Second, differentiating the inner energy aggregate gives:

$$\frac{\partial \ln e}{\partial e^e} = \frac{\alpha^e \left(e^e\right)^{-\frac{1}{\lambda}}}{\sum_{f \in \mathbb{F}} \alpha^f \left(e^f\right)^{\frac{\lambda-1}{\lambda}}}.$$

The electricity first-order condition is therefore:

$$p^e = \frac{\sigma - 1}{\sigma} R^* \alpha^E \frac{\alpha^e \left(e^e\right)^{-\frac{1}{\lambda}}}{\sum_{f \in \mathbb{F}} \alpha^f \left(e^f\right)^{\frac{\lambda-1}{\lambda}}}.$$

Multiplying by e^e yields:

$$R^* = \frac{\sigma}{\sigma - 1} \frac{p^e e^e}{\alpha^E} \left[1 + \sum_{f \neq e} \frac{\alpha^f}{\alpha^e} \left(\frac{e^f}{e^e}\right)^{\frac{\lambda-1}{\lambda}}\right].$$

Both Hicks-neutral productivity and the demand shifter have been eliminated using the electricity first-order condition. This derivation does not require output prices or physical output quantities.

For estimation, I write observed net revenue as $R = R^* \exp(u)$ where u is a log-revenue measurement disturbance satisfying $\mathbb{E}[u \mid \mathcal{I}] = 0$. The information set $\mathcal{I}$ includes input prices and the productivity and demand conditions known when inputs are chosen; the chosen inputs are measurable with respect to this information set. Substituting the calibrated energy exponent and defining $\widetilde{R} = \alpha^E R / (p^e e^e)$ gives:

$$\ln \widetilde{R} = \ln\left(\frac{\sigma}{\sigma - 1}\right) + \ln\left[1 + \sum_{f \neq e} \frac{\alpha^f}{\alpha^e} \left(\frac{e^f}{e^e}\right)^{\frac{\lambda-1}{\lambda}}\right] + u. \tag{A.16}$$

I estimate σ, λ, and the relative fuel weights by nonlinear least squares. Identification requires sufficient independent variation in relative fuel quantities and the stated conditional mean restriction. An unrestricted additional intercept would absorb the markup term and prevent separate identification of σ. The common scale of the fuel weights is

fixed by $\sum_{f\in\mathbb{F}} \alpha^f = 1$. Thus,

$$\alpha^e = \left[1 + \sum_{f\neq e} \left(\frac{\alpha^f}{\alpha^e}\right)\right]^{-1}, \qquad \alpha^f = \alpha^e \left(\frac{\alpha^f}{\alpha^e}\right), \quad f \neq e.$$

After estimation, the energy composite is constructed as:

$$e_t = \left[\sum_{f\in\mathbb{F}} \alpha^f \left(e_t^f\right)^{\frac{\lambda-1}{\lambda}}\right]^{\frac{\lambda}{\lambda-1}}. \tag{A.17}$$

F.2 Synthetic Firm Portfolios and Within-Firm Concentration

Statistical-disclosure restrictions permit the release of sector-level within-firm concentration moments and aggregate headquarters–plant–location masses, but not the firm–plant incidence matrix. Hence, I construct synthetic plant portfolios solely to approximate the concentration moment. They do not determine equilibrium price indices or replace the common plant-mass weights used in the equilibrium calculation. For each sector j and headquarters district i, the model represents firms by K_ϕ equally weighted productivity nodes ϕ_{ijk}. Following quantitative models of multinational production, I treat a firm's portfolio of production locations as a firm-level object shaped by firm heterogeneity (Tintelnot 2017; Arkolakis et al. 2023).

In the model's nested CES structure, holding sector, headquarters, and plant-location components fixed, the firm price index satisfies $P_{in,j}(\phi) = \widetilde{P}_{in,j}/\phi$ where $\widetilde{P}_{in,j} \equiv P_{in,j}(1)$ denotes the firm-level price index for a firm with unit productivity. Firm revenue and operating profit are therefore proportional to $P_{in,j}^{1-\sigma_j}(\phi)$, and hence to ϕ^{σ_j-1}. The released plant mass implies $K_\phi M_{il,j}/M_{i,j}$ plant slots at location l. I convert these targets into integer location totals using deterministic cumulative rounding and allocate the slots across productivity nodes in proportion to $\phi_{ijk}^{\sigma_j-1}$, using a largest-remainder rule so that each rounded location total is preserved.

Let n_{ijkl} denote the number of plants assigned to synthetic firm k at location l. For each assigned plant $h = 1, \ldots, n_{ijkl}$, I draw z_{ijklh} independently across plants and independently of ϕ_{ijk}, from the specified log-normal distribution. I evaluate $R_{il,j}\left(\phi_{ijk}, z_{ijklh}\right)$ at the equilibrium price indices, without feeding the sampled portfolio or its productivity draws back into those indices.

For portfolios with $\sum_l n_{ijkl} \geq 2$, I calculate:

$$s_{ijk}^{\max} = \frac{\max_{l,h:1\leq h\leq n_{ijkl}} R_{il,j}\left(\phi_{ijk}, z_{ijklh}\right)}{\sum_l \sum_{h=1}^{n_{ijkl}} R_{il,j}\left(\phi_{ijk}, z_{ijklh}\right)}.$$

I take the headquarters-mass-weighted mean across eligible synthetic firms, holding the simulated portfolios and productivity draws fixed across parameter evaluations. This procedure approximates the empirical concentration moment rather than an exact finite-

plant equilibrium. The portfolios match the released margins up to rounding, but the ownership links remain synthetic.

F.3 Matching in the Model

For plant-level moments, I construct model analogues of the empirical matched comparisons. Manufacturing plants are matched exactly by model sector and policy area, using Mahalanobis distance in baseline log sales and log electricity use. Power plants are matched by fuel type using baseline log sales, without exact area matching. I implement nearest-neighbor matching without replacement on the model's continuous plant masses. Each match uses equal treated and control mass, and matched mass cannot be reused. I average outcome differences using matched mass, both within and across sector or fuel groups. For equipment use, the empirical outcome is update-weighted log operating hours, whereas the model outcome is a change in continuous equipment intensity. I use the latter as a proxy on the empirical outcome's scale.

F.4 Estimation Criterion

I minimize a weighted sum of squared residuals. Empirical-response residuals are divided by their standard errors. Each receives weight one. For market k, the cap residual is $\sqrt{1+b_k^2}-1-b_k$, where $d_k<0$ is the cap-implied proportional emissions change, $\widehat{d_k}$ is its model counterpart, and $b_k=\left(d_k-\widehat{d_k}\right)/|d_k|$. This residual penalizes reductions exceeding the benchmark less than equally sized shortfalls.

G Solution Algorithm

I first describe the equilibrium solver conditional on emissions prices. I use the same production and electricity-market solver for the unregulated and policy equilibria. The unregulated equilibrium determines baseline wages, nominal incomes, and district-sector expenditures, without permit markets. In policy exercises, these variables and the capital rental rate remain fixed at their baseline values without endogenous labor market clearing and income updates, while plant choices, product prices, and the electricity price adjust. For cap-and-trade equilibria, an outer loop updates permit prices to satisfy the permit-market conditions. This simplification helps to focus on product, electricity, and pollution markets where the effect of the policy is studied.

1. Set permit-market caps and initialize permit prices.
2. Normalize the rental rate of capital to one: $r=1$.
3. Calculate the power plants' emission-control choices $a_{i,g}^0(z)$ (Equation (A.12)) and marginal generation cost schedules $\nu_{i,g}(y,z)$ (Equation (A.13)).

4. Guess an electricity price p^e and initialize wages w_i, sectoral price indices $P_{i,j}$, expenditures $X_{i,j}$, and owners' profits Π_i. Then, solve for the downstream manufacturing equilibrium.

   (a) Solve each manufacturing plant's emission-control problem, including boundary choices, to obtain the optimal equipment intensity $c_{il,j}(\phi, z)$ (Equation (A.8)), and calculate marginal production costs $\mu_{il,j}(\phi, z)$ (Equation (7)), input demands, sales, and profits $\pi_{il,j}(\phi, z)$ (Equation (A.11)).

   (b) Compute total manufacturing electricity demand ED, including electricity used to operate emission-control equipment (Equation (A.15)). Add the calibrated residual electricity demand ED_R to manufacturing demand ED, giving total electricity demand $ED + ED_R$. Dispatch generation subject to plant capacities $k_{i,g}(z)$. Given dispatched output, recover the output-dependent emission-control choice $a_{i,g}(y, z)$ (Equation (8)), plant profits $\pi_{i,g}(z)$ (Equation (A.14)), and the system marginal price.

   (c) Update wages w_i using the labor market clearing conditions (Equation (14)), sectoral price indices $P_{n,j}$ using the CES aggregators (Equation (10)), and expenditures $X_{i,j}$ using the product market clearing conditions (Equation (15)).

   (d) Repeat these updates until the manufacturing equilibrium converges.

5. Update the trial electricity price and repeat until the electricity price and SMP converge.

6. For cap-and-trade equilibria, update permit prices until the complementarity conditions (Equation (16)) hold.

7. Calculate real-income costs and avoided pollution damage.

H Additional Tables

Table A.1: Spatial distribution of policy areas

	SCA	CA	SA	SEA	Total
Number of regulated establishments	542	480	156	424	1,602
Contribution to $PM_{2.5}$ concentrations	21%	31%	10%	20%	82%
Population	49%	13%	5%	21%	88%
Land area	9%	14%	4%	11%	38%

Notes: SCA denotes the Seoul Capital Area, CA denotes the Central Area, SA denotes the Southern Area, and SEA denotes the South-East Area. The numbers of regulated establishments are calculated across all sectors from 2020 to 2024. Before linkage to the census, plant names and addresses serve as establishment identifiers in the policy registry. Concentration contributions refer to modeled contributions to domestic $PM_{2.5}$ concentrations, accounting for emissions and meteorological conditions. The population and land area are measured in 2018.
Sources: Ministry of Environment.

Table A.2: Policy goals

Policy area	Pollutant	Caps		Reduction rates
		Initial year (2020)	Final year (2024)	
SCA	Total	46,208	36,842	−20%
	NO_x	34,105	28,635	−16%
	SO_x	11,627	7,814	−33%
	TSP	476	393	−17%
CA	Total	160,616	108,094	−33%
	NO_x	114,678	82,452	−28%
	SO_x	43,145	23,611	−45%
	TSP	2,793	2,031	−27%
SA	Total	87,899	64,985	−26%
	NO_x	57,361	41,790	−27%
	SO_x	29,847	22,709	−24%
	TSP	691	486	−30%
SEA	Total	135,629	91,999	−32%
	NO_x	90,741	58,956	−35%
	SO_x	43,534	32,056	−26%
	TSP	1,354	987	−27%

Notes: SCA denotes the Seoul Capital Area, CA denotes the Central Area, SA denotes the Southern Area, and SEA denotes the South-East Area. Caps are the sums of emission permits (initial allocation (including additional allocation), carryover, and net purchases (purchases – sales)) of all regulated establishments within markets. Total is the sum of NO_x caps, SO_x caps, and TSP caps (tons). Reduction rates are taken between 2020 and 2024 and expressed as percentages.
Sources: Ministry of Environment.

Table A.3: Summary statistics: sectoral distribution

	Counts (%)		Emissions (%)		
Sector	Plants	Firms	NO_x	SO_x	TSP
Manufacturing					
Mining/quarrying	0.1	0.1	0.04	0.03	0.004
Food/beverages/tobacco	9.2	7.8	0.3	0.002	0.004
Textiles/apparel/leather	4.1	5.2	0.9	1.3	6.0
Wood/paper/printing/media	6.6	7.5	0.2	0.1	0.2
Petro/chem/pharma/rubber/plastics	19.3	19.4	12.1	35.5	13.5
Non-metallic minerals	14.0	14.8	14.7	7.2	11.4
Basic/fabricated metals	18.0	18.7	20.1	20.9	35.1
Electronics/instruments/electrical	7.6	7.2	0.6	0.003	0.01
Machinery/transport (incl. auto)	5.8	6.2	1.1	0.2	0.03
Furniture/misc./repair	0.6	0.6	2.7	3.3	16.5
Power generation					
Fossil-fuel power generation	13.7	11.3	35.7	29.6	15.0
Non-fossil-fuel power generation	1.1	1.4	1.7	0.5	0.4
Total	100	100	90.1	98.5	98.1

Notes: The first two columns use the 2017 KSIC classifications and report the percentage of regulated plants and firms. The last three columns use 2017 emissions data and report each sector's share of manufacturing and power generation emissions in regulated districts by pollutant. The percentages do not sum to 100% because some categories that overlap across sectors are dropped.
Sources: Ministry of Environment, Statistics Korea.

Table A.4: Summary statistics: district–industry

	Manufacturing		*Power generation*	
	Targeted	Never	Targeted	Never
ln NO_x emissions	6.17	5.61	10.62	11.25
	(3.22)	(2.94)	(3.28)	(3.12)
ln SO_x emissions	1.02	1.02	6.69	8.48
	(4.20)	(4.35)	(4.79)	(4.22)
ln TSP emissions	−0.25	−0.35	6.75	7.40
	(3.96)	(3.88)	(3.70)	(3.49)
N	3,641	1,119	127	17

Notes: Observations are district–industry cells in 2017. Emissions are measured in kilograms. Targeted (Never) denotes cells in regulated (unregulated) districts. Standard deviations are shown in parentheses.
Sources: Ministry of Environment.

Table A.5: Summary statistics: plant

Manufacturing	Pre-matching			Post-matching		
Regulated	Targeted	Never	StdDif	Targeted	Never	StdDif
1{Equipment update}	0.16	0.11	0.13	0.16	0.95	−2.66
	(0.37)	(0.32)		(0.36)	(0.22)	
ln Equipment use × 1{Equipment update}	1.37	0.91	0.16	1.36	7.78	−2.50
	(3.17)	(2.52)		(3.16)	(1.86)	
ln Sales	11.61	7.69	2.43	11.45	9.82	1.16
	(1.71)	(1.51)		(1.51)	(1.28)	
ln Number of employees	4.82	2.20	2.02	4.72	3.67	0.91
	(1.43)	(1.16)		(1.29)	(0.97)	
ln Electricity use	7.44	2.64	2.41	7.50	5.18	1.34
	(2.22)	(1.73)		(1.86)	(1.49)	
ln Electricity share	−1.08	−0.46	−0.59	−1.07	−0.35	−0.86
	(1.06)	(1.05)		(0.87)	(0.58)	
N	701	85,231		409		

Manufacturing	Pre-matching			Post-matching		
Unregulated affiliated	Targeted	Never	StdDif	Targeted	Never	StdDif
1{Equipment update}	0.21	0.11	0.27	0.28	0.95	−1.90
	(0.41)	(0.32)		(0.45)	(0.22)	
ln Equipment use × 1{Equipment update}	1.79	0.91	0.29	2.33	7.72	−1.85
	(3.46)	(2.52)		(3.75)	(1.94)	
ln Sales	10.64	7.69	1.75	10.94	10.08	0.66
	(1.85)	(1.51)		(1.46)	(1.07)	
ln Number of employees	4.00	2.20	1.32	4.31	3.74	0.54
	(1.54)	(1.16)		(1.23)	(0.92)	
ln Electricity use	5.85	2.64	1.51	6.36	5.26	0.61
	(2.46)	(1.73)		(2.07)	(1.39)	
ln Electricity share	−0.53	−0.46	−0.08	−0.64	−0.33	−0.43
	(0.82)	(1.05)		(0.68)	(0.60)	
N	562	85,231		96		

Table A.5 (continued): Power generation

Power generation	Pre-matching			Post-matching		
Regulated	Targeted	Never	StdDif	Targeted	Never	StdDif
1{Equipment update}	0.05	0.04	0.01	0.07	0.71	−3.29
	(0.21)	(0.20)		(0.27)	(0.47)	
ln Equipment use × 1{Equipment update}	0.35	0.32	0.02	0.65	5.70	−3.24
	(1.65)	(1.53)		(2.43)	(3.85)	
ln Sales	10.70	4.07	3.03	9.89	8.65	0.73
	(2.43)	(1.91)		(2.27)	(3.47)	
ln Number of employees	3.80	1.04	2.17	2.93	3.18	−0.20
	(1.45)	(1.07)		(1.87)	(1.97)	
N	131	74,633		14		

Notes: Electricity and fuel use are measured in tons of oil equivalent, sales are measured in 1,000 U.S. dollars, employment is calculated as the number of permanent employees, and the use of emission-control equipment is measured in hours per year. Equipment update indicates whether a plant updates its emission-control equipment. Since not every manufacturing plant uses fossil fuels, I use ln Electricity share, computed as ln (Electricity/(Electricity+Coal+Oil+Gas)). Exchange rates are sourced from the Federal Reserve Bank of St. Louis. StdDif denotes the standardized mean difference. Post-matching differences are normalized by the pooled treated and control standard deviations in the matching-eligible sample before matching. For manufacturing plants, I focus on those that have electricity use. Targeted denotes the treated group in each panel and Never the never-treated comparison plants. Post-matching *N* is the number of matched pairs. Standard deviations are shown in parentheses.
Sources: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea, Federal Reserve Bank of St. Louis.

Table A.6: Placebo: district–industry changes

	Manufacturing	*Power generation*
ln NO_x emissions	−0.44***	−0.15
	(0.05)	(0.17)
t-statistic	−8.59	−0.90
RI-t p-value	0.000	0.443
ln SO_x emissions	−0.46***	−0.41
	(0.08)	(0.25)
t-statistic	−5.74	−1.64
RI-t p-value	0.002	0.092
ln TSP emissions	−0.40***	−0.09
	(0.06)	(0.21)
t-statistic	−6.30	−0.44
RI-t p-value	0.002	0.723

Notes: Within each sector, I draw 499 distinct placebo assignments by randomly reassigning treatment across districts, holding fixed the number of treated districts and excluding the actual assignment. All district–industry cells within a district receive the same treatment status. Actual and placebo regressions use the same fixed effects and district–industry-clustered standard errors (Equation (2)). The reported randomization-inference p-value based on the t-statistic (RI-t p-value) is the fraction of valid placebo assignments for which the absolute placebo t-statistic exceeds the absolute actual t-statistic. Given the nonrandom geographic targeting of regulation, I interpret these comparisons as supplementary placebo evidence rather than exact design-based inference. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment.

Table A.7: District–industry changes, excluding the Seoul Capital Area

	Manufacturing	*Power generation*
ln NO_x emissions	−0.40***	−0.32
	(0.05)	(0.23)
ln SO_x emissions	−0.43***	−0.10
	(0.08)	(0.24)
ln TSP emissions	−0.36***	−0.41*
	(0.06)	(0.19)
N	25,896	652

Notes: The average effects are calculated from Equation (2) without the observations from the Seoul Capital Area. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment.

Table A.8: Geographic pollution shifts, excluding the Seoul Capital Area

	Manufacturing	*Power generation*
ln NO_x emissions	0.32	0.53
	(0.18)	(0.61)
ln SO_x emissions	0.77*	1.44*
	(0.38)	(0.65)
ln TSP emissions	0.38	−0.23
	(0.24)	(0.95)
N	2,214	101

Notes: The average effects are calculated from Equation (4) without the observations from the Seoul Capital Area. Both the sample and the distance-based exposure measure are recomputed after excluding the Seoul Capital Area. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment.

Table A.9: District–industry changes, controlling for baseline-emissions trends

	Manufacturing	*Power generation*
ln NO_x emissions	−0.20***	0.19
	(0.05)	(0.19)
ln SO_x emissions	−0.02	−0.17
	(0.08)	(0.53)
ln TSP emissions	−0.12*	0.02
	(0.06)	(0.12)
N	38,419	1,143

Notes: The average effects are calculated from Equation (2). The specification additionally includes interactions between mean ln Emissions over 2016–2019 and year indicators. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment.

Table A.10: Geographic pollution shifts, controlling for baseline-emissions trends

	Manufacturing	*Power generation*
ln NO_x emissions	0.36*	−0.19
	(0.18)	(0.71)
ln SO_x emissions	0.88*	0.93
	(0.35)	(0.77)
ln TSP emissions	0.46	−0.78
	(0.24)	(1.08)
N	1,979	86

Notes: The average effects are calculated from Equation (4). The specification additionally includes interactions between mean ln Emissions over 2016–2019 and year indicators. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment.

Table A.11: District–industry changes, excluding Ongjin

	Manufacturing	*Power generation*
ln NO_x emissions	−0.44***	−0.15
	(0.05)	(0.17)
ln SO_x emissions	−0.46***	−0.41
	(0.08)	(0.25)
ln TSP emissions	−0.40***	−0.09
	(0.06)	(0.21)
N	39,068	1,209

Notes: The average effects are calculated from Equation (2). These estimates exclude Ongjin district, which is partly regulated. Since Ongjin district is already excluded from the geographic pollution shift samples for both manufacturing and power generation, this restriction leaves the geographic pollution shift estimates unchanged. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment.

Table A.12: District–industry changes, with district clusters

	Manufacturing	*Power generation*
ln NO_x emissions	−0.44***	−0.15
	(0.12)	(0.14)
ln SO_x emissions	−0.46***	−0.41
	(0.12)	(0.23)
ln TSP emissions	−0.40***	−0.09
	(0.11)	(0.20)
N	39,086	1,219

Notes: The average effects are calculated from Equation (2). Standard errors are clustered at the district level; the specification and sample otherwise match the baseline. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment.

Table A.13: Geographic pollution shifts, with district clusters

	Manufacturing	*Power generation*
ln NO_x emissions	0.35	0.07
	(0.29)	(0.47)
ln SO_x emissions	1.42*	1.05*
	(0.58)	(0.46)
ln TSP emissions	0.61	−0.52
	(0.40)	(0.36)
N	2,105	101

Notes: The average effects are calculated from Equation (4). Standard errors are clustered at the district level; the specification and sample otherwise match the baseline. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment.

Table A.14: District–industry changes, with a fully balanced sample

	Manufacturing	*Power generation*
ln NO_x emissions	−0.32***	−0.19
	(0.05)	(0.16)
ln SO_x emissions	−0.36***	−0.48
	(0.09)	(0.29)
ln TSP emissions	−0.28***	−0.13
	(0.07)	(0.20)
N	32,216	976

Notes: The average effects are calculated from Equation (2). The sample retains district–industry cells with positive, non-missing emissions of all three pollutants in every year. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment.

Table A.15: Geographic pollution shifts, with a fully balanced sample

	Manufacturing	*Power generation*
ln NO_x emissions	−0.05	0.54
	(0.21)	(0.63)
ln SO_x emissions	1.17*	1.44
	(0.51)	(0.70)
ln TSP emissions	0.22	−0.22
	(0.30)	(0.98)
N	1,392	72

Notes: The average effects are calculated from Equation (4). The sample retains district–industry cells with positive, non-missing emissions of all three pollutants in every year. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment.

Table A.16: Plant responses, GHG ETS plants

	Manufacturing		*Power generation*
	Regulated	Unregulated affiliated	Regulated
1{Equipment update}	0.03*	−0.01	0.20
	(0.01)	(0.03)	(0.11)
ln Equipment use × 1{Equipment update}	0.23*	−0.07	1.53
	(0.11)	(0.24)	(0.86)
ln Sales	0.07	0.17	0.14
	(0.04)	(0.09)	(0.46)
ln Electricity use	1.12***	0.68*	
	(0.10)	(0.26)	
ln Electricity share	0.10**	0.09	
	(0.04)	(0.09)	
N	3,920	848	160

Notes: The average effects are calculated from Equation (6). The sample includes plants covered by the greenhouse gas emissions trading system at least once. Regulated and unregulated-affiliate status in this table refer to the local-air-pollution policy, not the greenhouse gas emissions trading system. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea.

Table A.17: Plant responses, GHG ETS and pilot plants

	Manufacturing		*Power generation*
	Regulated	Unregulated affiliated	Regulated
1{Equipment update}	0.03*	0.04	0.20
	(0.01)	(0.03)	(0.11)
ln Equipment use × 1{Equipment update}	0.22*	0.29	1.53
	(0.10)	(0.23)	(0.86)
ln Sales	0.09**	0.23**	0.40
	(0.03)	(0.07)	(0.51)
ln Electricity use	1.15***	0.61**	
	(0.09)	(0.20)	
ln Electricity share	0.11**	0.07	
	(0.03)	(0.07)	
N	4,608	1,200	160

Notes: The average effects are calculated from Equation (6). The sample includes plants covered by the greenhouse gas emissions trading system or its pilot program at least once. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea.

Table A.18: Plant responses, excluding the Seoul Capital Area

	Manufacturing		*Power generation*
	Regulated	Unregulated affiliated	Regulated
1{Equipment update}	0.02	0.05	0.15
	(0.01)	(0.03)	(0.14)
ln Equipment use × 1{Equipment update}	0.14	0.46*	0.99
	(0.11)	(0.22)	(1.21)
ln Sales	0.06	0.17*	0.53
	(0.04)	(0.07)	(0.45)
ln Electricity use	1.07***	0.32	
	(0.10)	(0.23)	
ln Electricity share	0.09*	0.01	
	(0.04)	(0.06)	
N	4,528	1,200	160

Notes: The average effects are calculated from Equation (6) without plants in the Seoul Capital Area. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea.

Table A.19: Plant responses, without exact matching on policy areas

	Manufacturing		*Power generation*
	Regulated	Unregulated affiliated	Regulated
1{Equipment update}	0.02*	0.03	0.13
	(0.01)	(0.02)	(0.11)
ln Equipment use × 1{Equipment update}	0.18*	0.24	0.90
	(0.08)	(0.18)	(0.89)
ln Sales	0.06*	0.12*	0.53
	(0.03)	(0.05)	(0.38)
ln Electricity use	0.75***	0.08	
	(0.07)	(0.16)	
ln Electricity share	0.08*	0.01	
	(0.03)	(0.05)	
N	7,072	1,648	224

Notes: The average effects are calculated from Equation (6) without an exact match on policy areas. Electricity outcomes are not reported for power plants because their electricity use is not observed. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea.

Table A.20: Plant responses, no same-industry affiliates

	Manufacturing	
	Regulated	Unregulated affiliated
1{Equipment update}	0.03	
	(0.02)	
ln Equipment use × 1{Equipment update}	0.19	
	(0.12)	
ln Sales	0.07	−0.08
	(0.04)	(0.16)
ln Electricity use	0.91***	1.85**
	(0.12)	(0.41)
ln Electricity share	0.07	0.78
	(0.05)	(0.51)
N	3,568	96

Notes: The average effects are calculated from Equation (6). The sample for the first column keeps regulated manufacturing plants without any unregulated affiliated plants in the same three-digit KSIC industry. The sample for the second column keeps unregulated affiliated plants that do not share the same three-digit KSIC industry as the corresponding regulated counterparts. I do not consider a counterpart for power plants because within-firm leakage is unlikely to arise. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea.

Table A.21: Plant responses, affiliates outside vs. inside regulated districts

	Manufacturing	
Unregulated affiliated	Outside	Inside
1{Equipment update}	−0.03	0.08*
	(0.03)	(0.03)
ln Equipment use × 1{Equipment update}	−0.13	0.61*
	(0.28)	(0.26)
ln Sales	0.23*	0.17*
	(0.09)	(0.07)
ln Electricity use	0.18	0.45*
	(0.45)	(0.19)
ln Electricity share	0.14	0.08
	(0.10)	(0.08)
N	448	1,088

Notes: The average effects are calculated from Equation (6). The sample for the first column keeps unregulated affiliated manufacturing plants outside regulated districts while the sample for the second column keeps unregulated affiliated manufacturing plants inside regulated districts. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea.

Table A.22: Plant responses, firm ownership

	Power generation	
	Regulated	Unregulated affiliated
1{Equipment update}	0.12	0.08*
	(0.11)	(0.03)
ln Equipment use × 1{Equipment update}	0.80	0.63*
	(0.93)	(0.26)
ln Sales	0.22	−0.26
	(0.35)	(0.13)
N	208	1,056

Notes: The average effects are calculated from Equation (6). The sample applies the manufacturing ownership design to power plants: regulated power plants and unregulated power plants affiliated with regulated firms are each compared with matched unregulated unaffiliated power plants. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea.

Table A.23: Plant responses, propensity-score matching

	Manufacturing		*Power generation*
	Regulated	Unregulated affiliated	Regulated
1{Equipment update}	0.02*	0.03	0.13
	(0.01)	(0.02)	(0.11)
ln Equipment use × 1{Equipment update}	0.17*	0.29	0.90
	(0.08)	(0.17)	(0.89)
ln Sales	0.08*	0.15*	0.49
	(0.03)	(0.06)	(0.38)
ln Electricity use	0.94***	0.29	
	(0.08)	(0.21)	
ln Electricity share	0.10*	0.16	
	(0.04)	(0.08)	
N	6,544	1,536	224

Notes: The average effects are calculated from Equation (6) with propensity score matching. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea.

Table A.24: Plant responses, without the environmental-comparability restriction

	Manufacturing		*Power generation*
	Regulated	Unregulated affiliated	Regulated
1{Equipment update}	0.01	0.02	−0.01
	(0.01)	(0.02)	(0.01)
ln Equipment use × 1{Equipment update}	0.04	0.18	−0.06
	(0.06)	(0.15)	(0.07)
ln Sales	0.11***	0.11	0.46***
	(0.03)	(0.06)	(0.13)
ln Electricity use	0.65***	0.02	
	(0.07)	(0.16)	
ln Electricity share	0.09**	0.10	
	(0.03)	(0.06)	
N	7,232	1,632	1,616

Notes: The average effects are calculated from Equation (6). The control groups retain plants excluded by the baseline environmental-comparability restriction (Appendix C.1). Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea.

Table A.25: Plant responses, trimmed affiliate samples

	Manufacturing
	Unregulated affiliated
1{Equipment update}	0.02
	(0.05)
ln Equipment use × 1{Equipment update}	0.30
	(0.38)
ln Sales	0.33*
	(0.12)
ln Electricity use	0.10
	(0.41)
ln Electricity share	0.07
	(0.15)
N	464

Notes: The average effects are calculated from Equation (6). For the unregulated-affiliate comparison, I trim both the treatment and control groups by removing plants that are not environmentally comparable (Appendix C.1). Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea.

Table A.26: Plant responses, by number of regulated pollutants

	Manufacturing		
	Regulated		
Number of regulated pollutants	1	2	3
1{Equipment update}	0.02	0.02	0.05*
	(0.01)	(0.01)	(0.02)
ln Equipment use × 1{Equipment update}	0.16	0.15	0.41*
	(0.10)	(0.11)	(0.18)
ln Sales	0.01	0.09**	0.07
	(0.04)	(0.03)	(0.07)
ln Electricity use	0.97***	0.88***	1.13***
	(0.09)	(0.10)	(0.10)
ln Electricity share	0.14***	0.03	0.15*
	(0.04)	(0.05)	(0.06)
N	5,032	4,424	3,632

Notes: The average effects are calculated from Equation (6). Each column compares manufacturing plants regulated for the indicated number of pollutants with the full baseline matched control sample. Matching is conducted for the pooled sample, not separately by subgroup. Standard errors are shown in parentheses. * denotes p<0.05, ** denotes p<0.01, and *** denotes p<0.001.
Sources: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea.

Table A.27: Factor shares and CES fuel weights

Manufacturing	Factors			Fuels			
All plants	α^E	α^K	α^L	α^{Coal}	$\alpha^{\text{Electricity}}$	α^{Gas}	α^{Oil}
Mining/quarrying	0.29	0.06	0.65	0.11	0.63	<0.01	0.26
Food/beverages/tobacco	0.21	0.05	0.74	<0.01	0.64	0.17	0.19
Textiles/apparel/leather	0.20	0.04	0.77	<0.01	0.71	0.15	0.14
Wood/paper/printing/media	0.27	0.05	0.69	<0.01	0.39	0.31	0.31
Petro/chem/pharma/rubber/plastics	0.41	0.05	0.53	<0.01	0.48	0.25	0.27
Non-metallic minerals	0.38	0.06	0.56	<0.01	0.56	0.19	0.25
Basic/fabricated metals	0.33	0.06	0.61	<0.01	0.53	0.20	0.28
Electronics/instruments/electrical	0.18	0.06	0.76	0.19	0.70	0.04	0.07
Machinery/transport (incl. auto)	0.09	0.05	0.86	0.65	0.33	0.01	0.02
Furniture/misc./repair	0.06	0.03	0.91	<0.01	0.37	0.32	0.31

Notes: The table reports sector-specific Cobb–Douglas factor shares and CES fuel weights estimated using the baseline sample of all plants.
Sources: MMS and SBR.

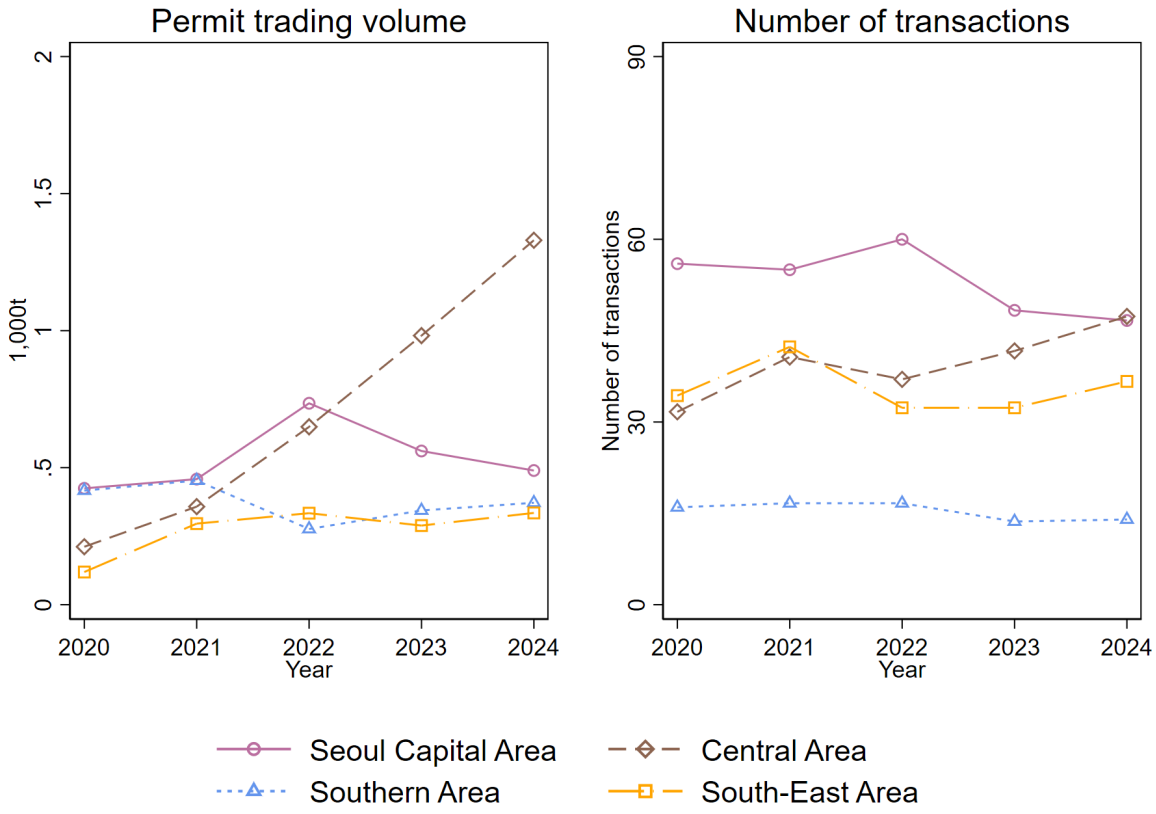
Table A.28: Changes across policy areas

Changes / Policy area	Consumption change (%)	Pollution damage change (%)
Seoul Capital Area	−1.39	−35.0
Central Area	−1.46	−35.8
Southern Area	−1.57	−29.5
South-East Area	−1.44	−47.4

Notes: Entries are percentage changes relative to the unregulated equilibrium. Each area includes only districts covered by the actual policy. Pollution damage accounts for emissions from all manufacturing and power plants in those districts, regardless of plant-level eligibility. Consumption and pollution damage are summed within each area before computing percentage changes.

I Additional Figures

Figure A.3: Permit trading



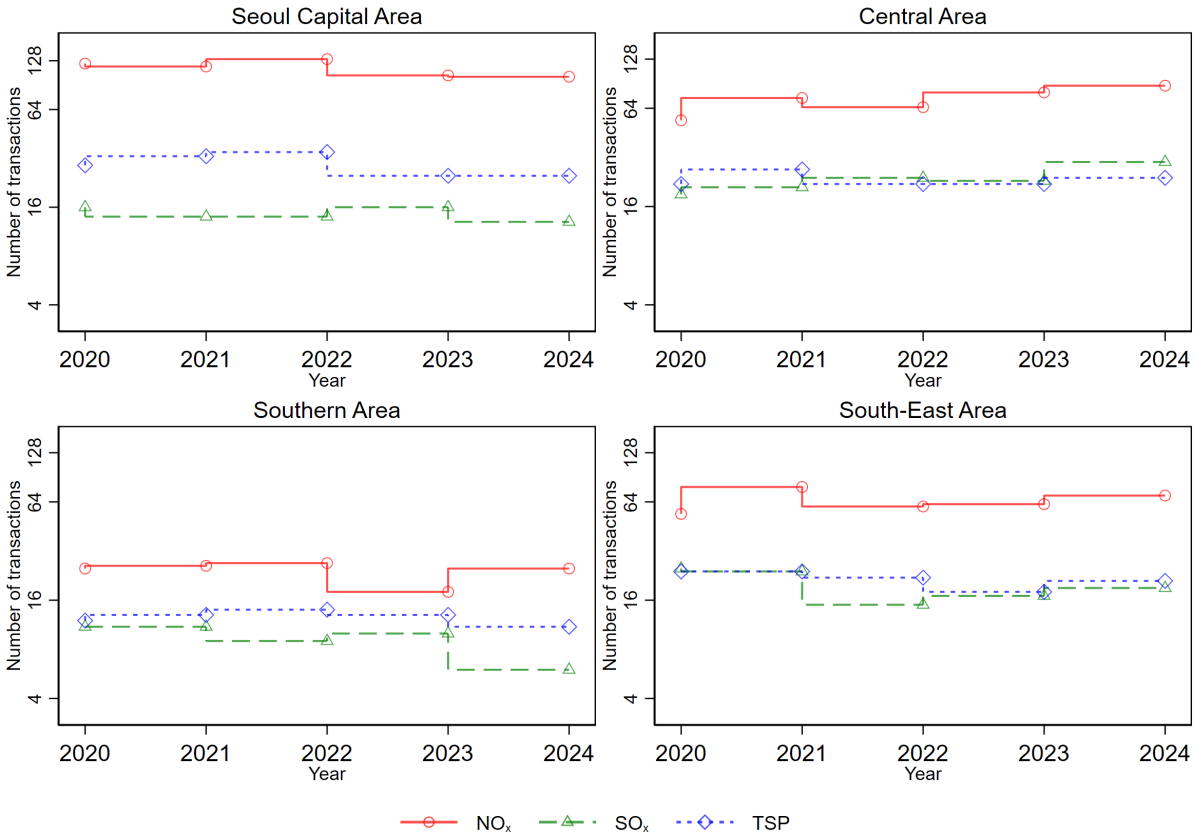
Notes: The panels report trading volumes and transaction counts averaged across markets, including transactions without consideration. Low trading volumes need not indicate trading frictions: firms may require few permit purchases when abatement is inexpensive or initial allocations are generous. Declining permit prices and increased abatement are consistent with limited demand for additional permits, although these patterns do not rule out trading frictions. These records also show active permit trading, with most recorded purchase requests fulfilled but substantial posted supply remaining unsold.
Sources: Ministry of Environment.

Figure A.4: Permit trading volume by pollutant



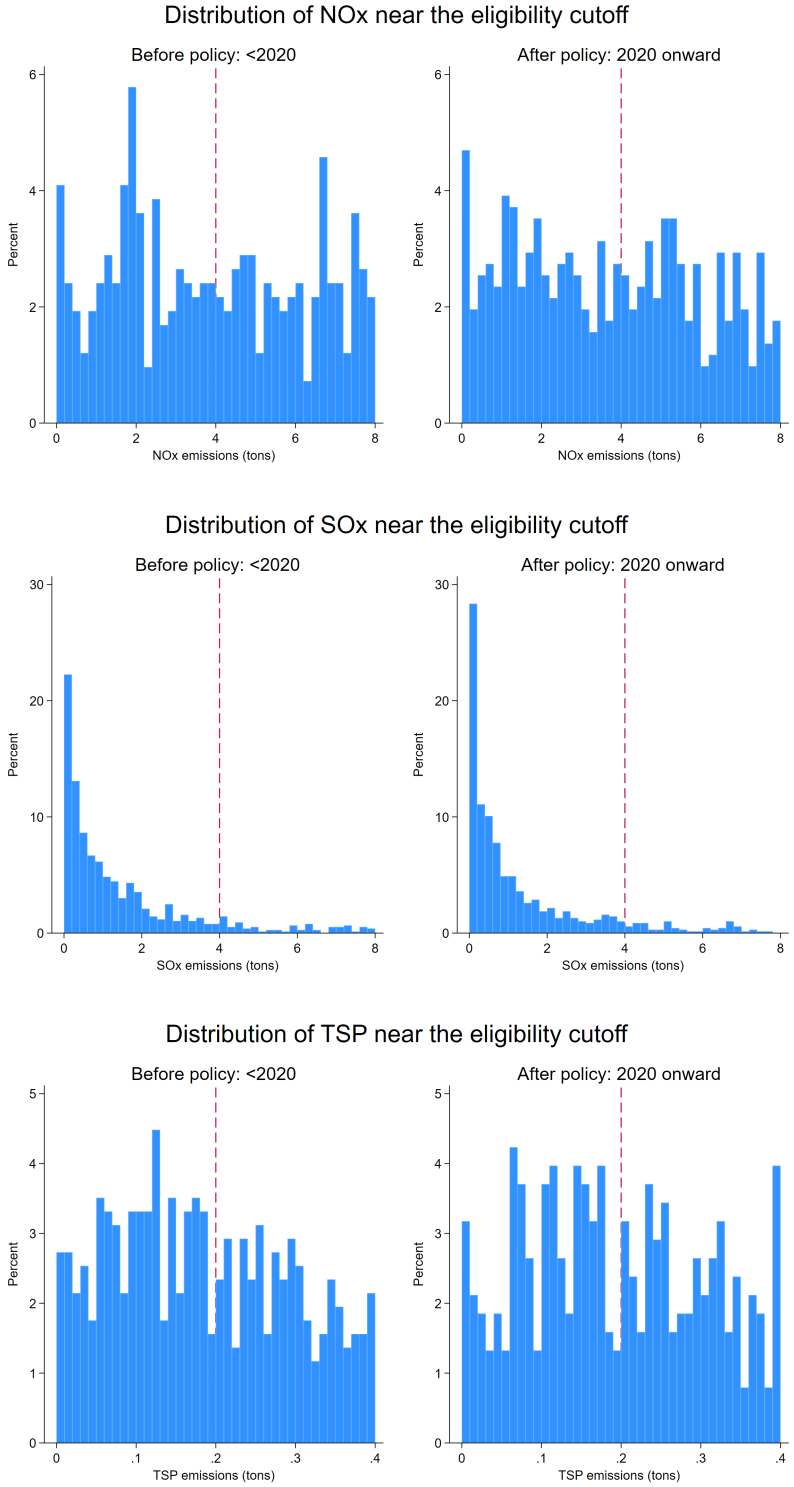
Notes: The volume includes transactions without consideration.
Sources: Ministry of Environment.

Figure A.5: Number of transactions by pollutant



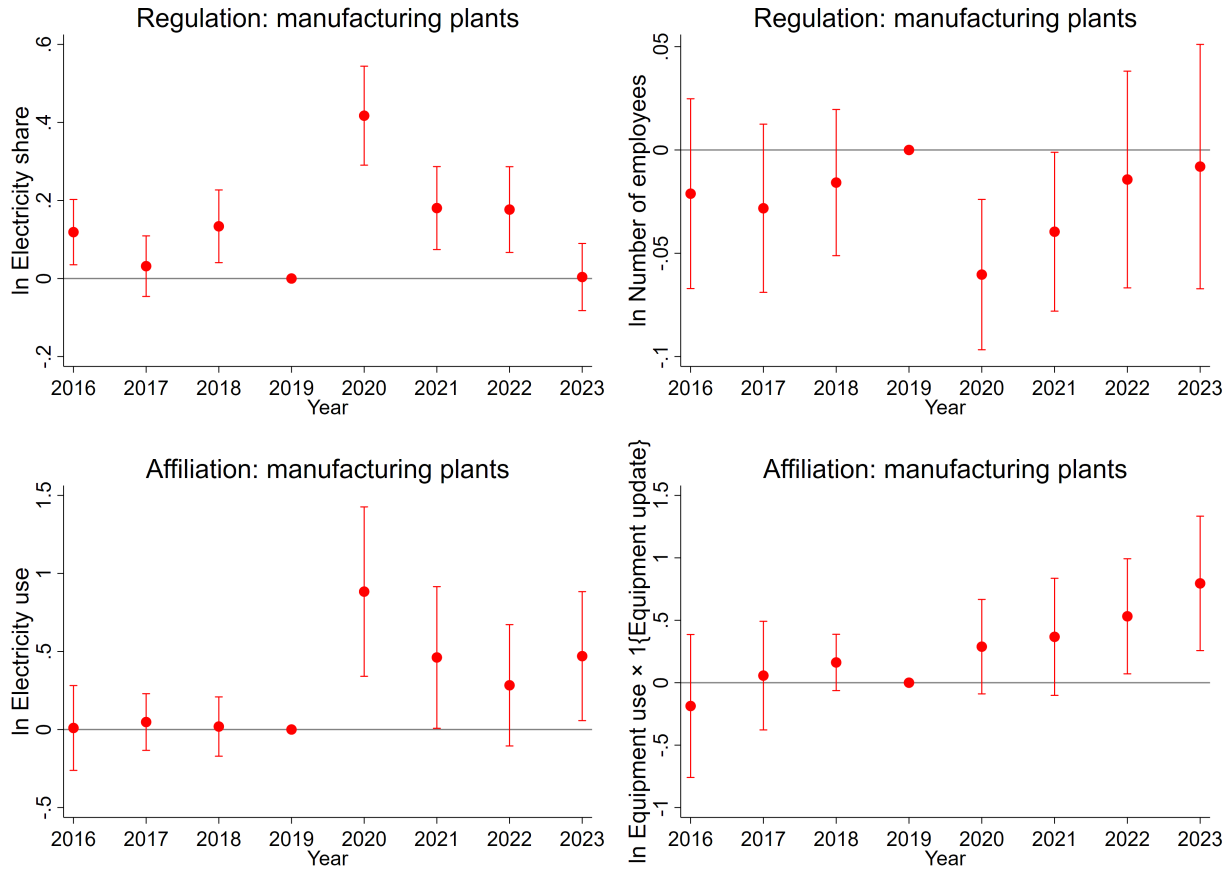
Notes: The number of transactions includes transactions without consideration.
Sources: Ministry of Environment.

Figure A.6: Plant-level emissions distributions



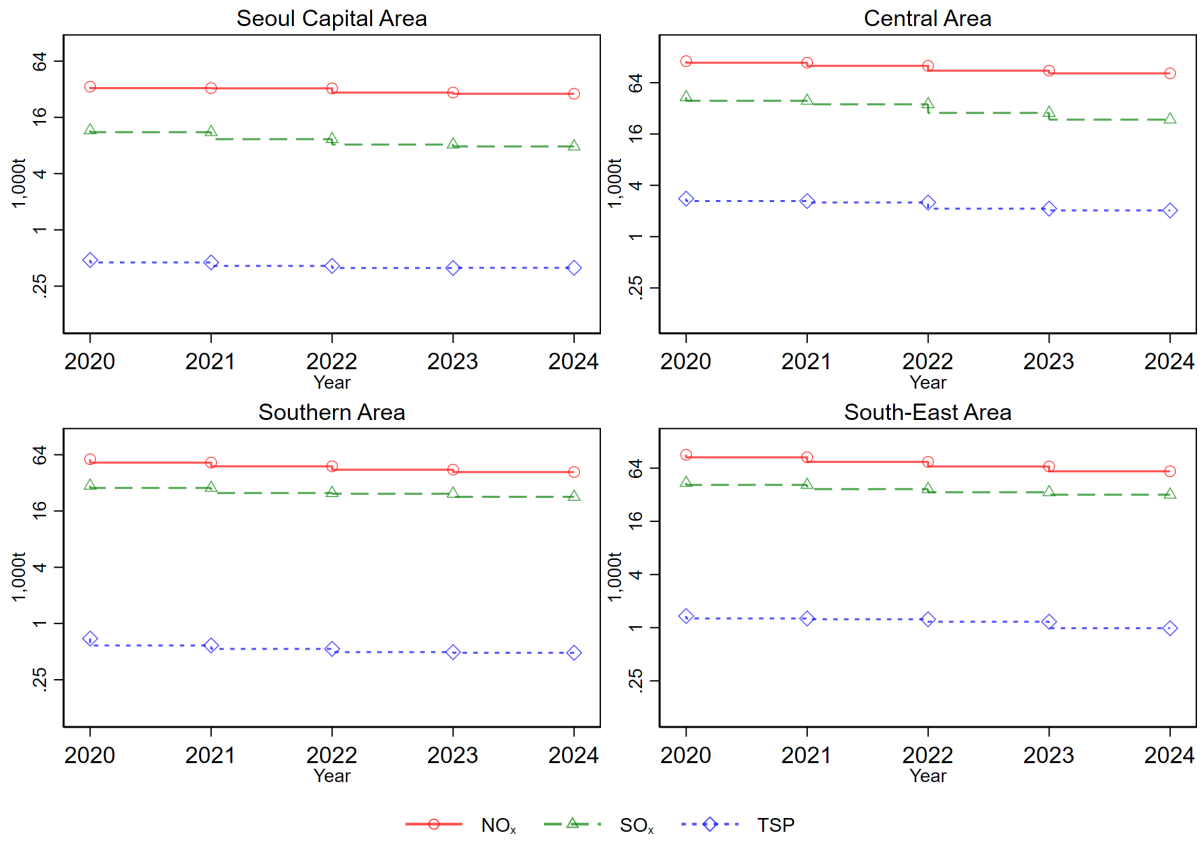
Notes: These panels show positive plant-level emissions in regulated districts within twice the pollutant-specific cutoff. Bin widths are 0.2 tons for NO_x and SO_x and 0.01 tons for TSP. Dashed lines mark the cutoffs. Percentages are calculated separately within each displayed sample.
Sources: Ministry of Environment.

Figure A.7: Event dynamics: additional plant outcomes



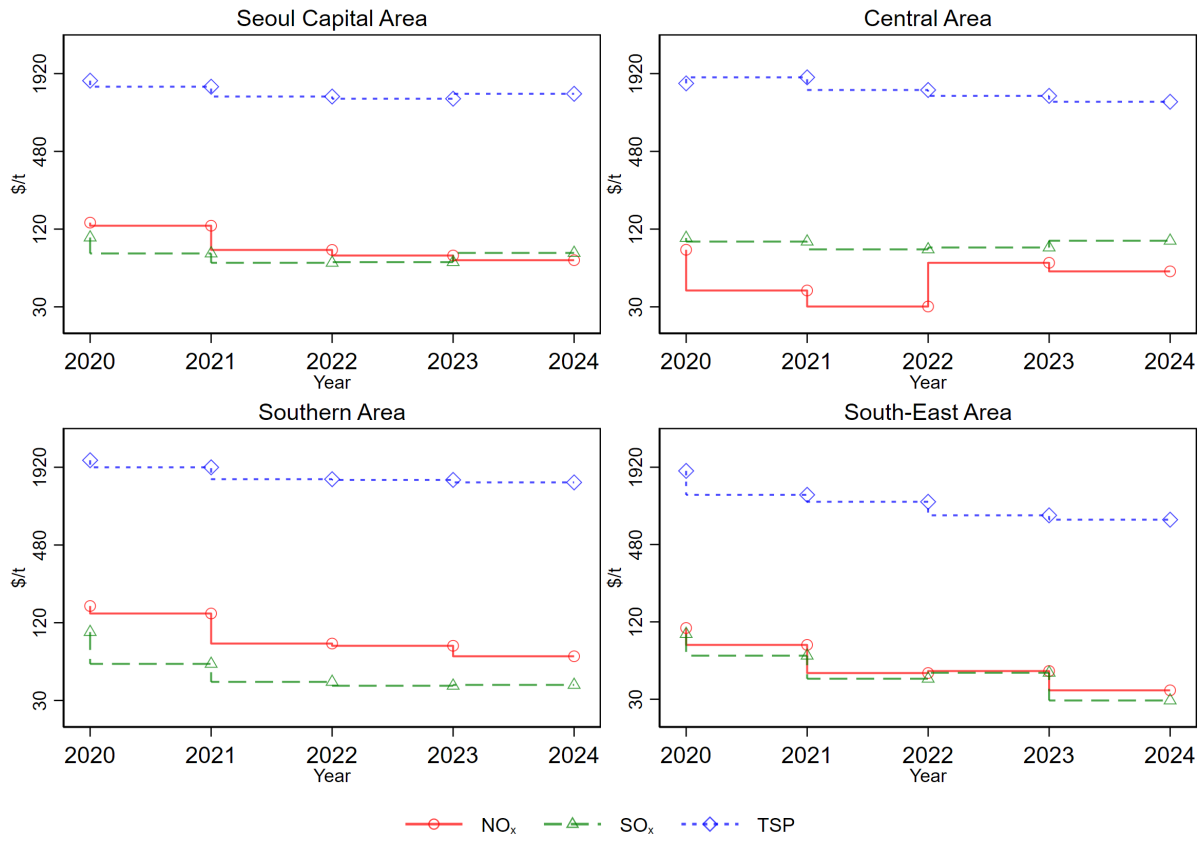
Notes: The upper panels report relative changes in electricity shares and employment at regulated manufacturing plants. The lower panels report relative changes in electricity use and update-weighted log equipment use at unregulated manufacturing affiliates. The results are calculated from Equation (5). The 95% confidence intervals are shown jointly with the point estimates.
Sources: Ministry of Environment, Ministry of the Interior and Safety, Statistics Korea.

Figure A.8: Caps by pollutant



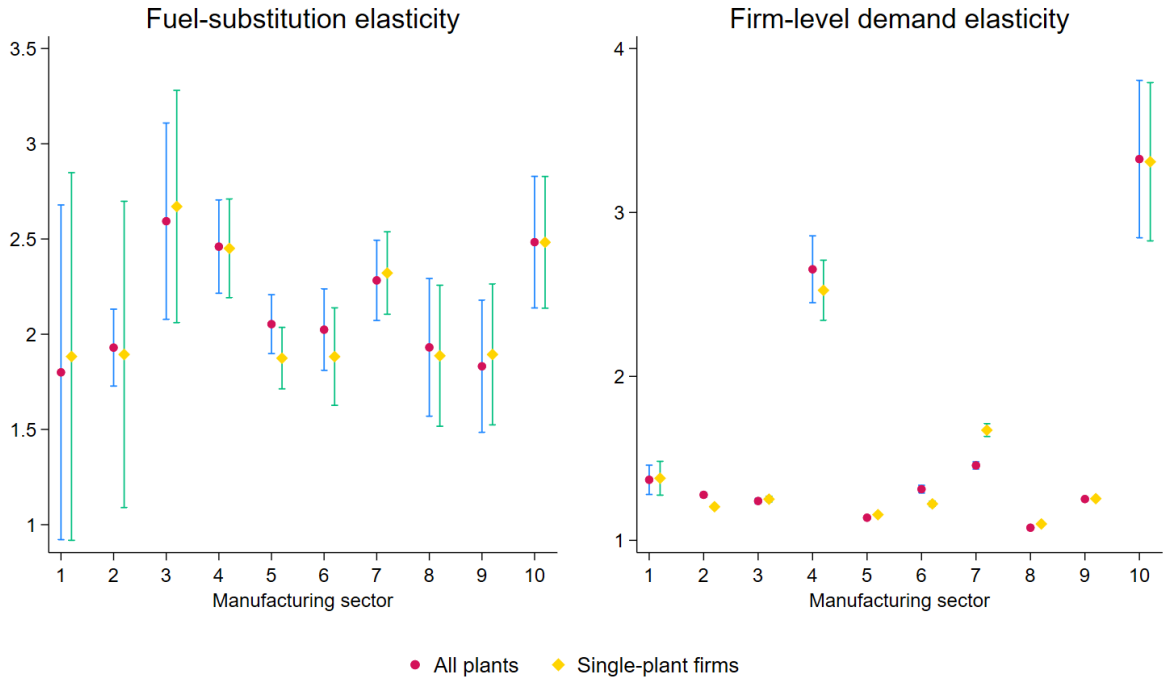
Notes: These panels plot the total amount of permits in each market, which is defined by a pollutant and a policy area. Across all markets, there is a downward trend over time.
Sources: Ministry of Environment.

Figure A.9: Average permit prices by pollutant



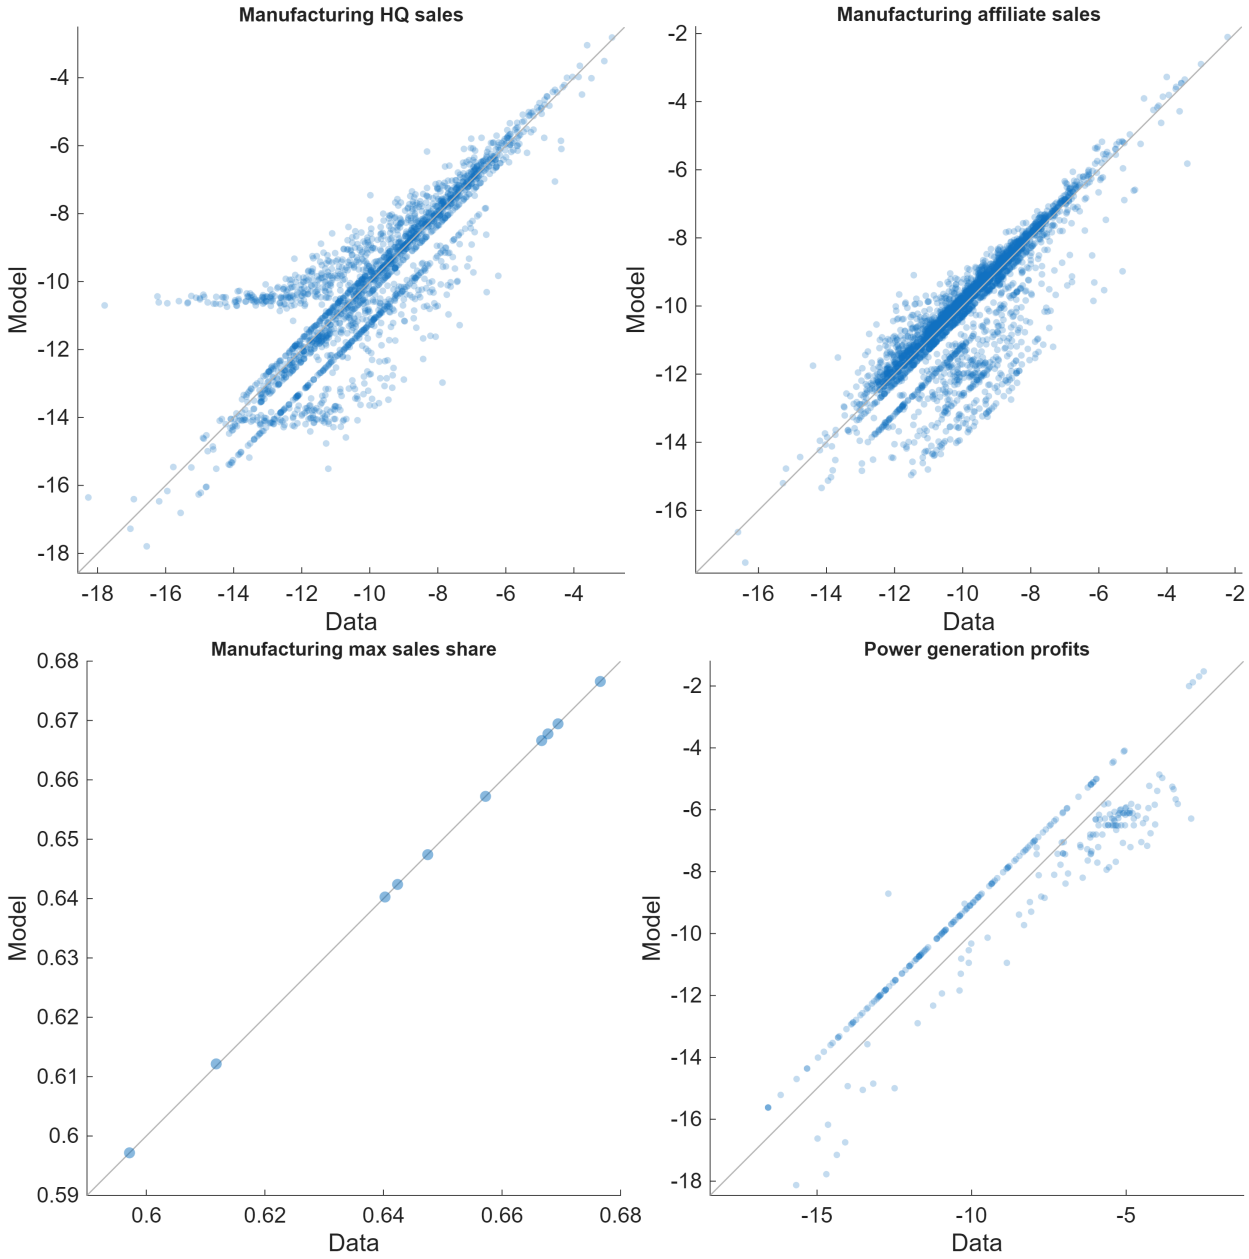
Notes: The average transaction unit price is calculated as the annual transaction amount divided by the annual transaction volume in each area, excluding transactions without consideration. Here I use annual average spot exchange rates to convert prices from Korean won to U.S. dollars.
Sources: Ministry of Environment, Federal Reserve Bank of St. Louis.

Figure A.10: Fuel-substitution and demand elasticities



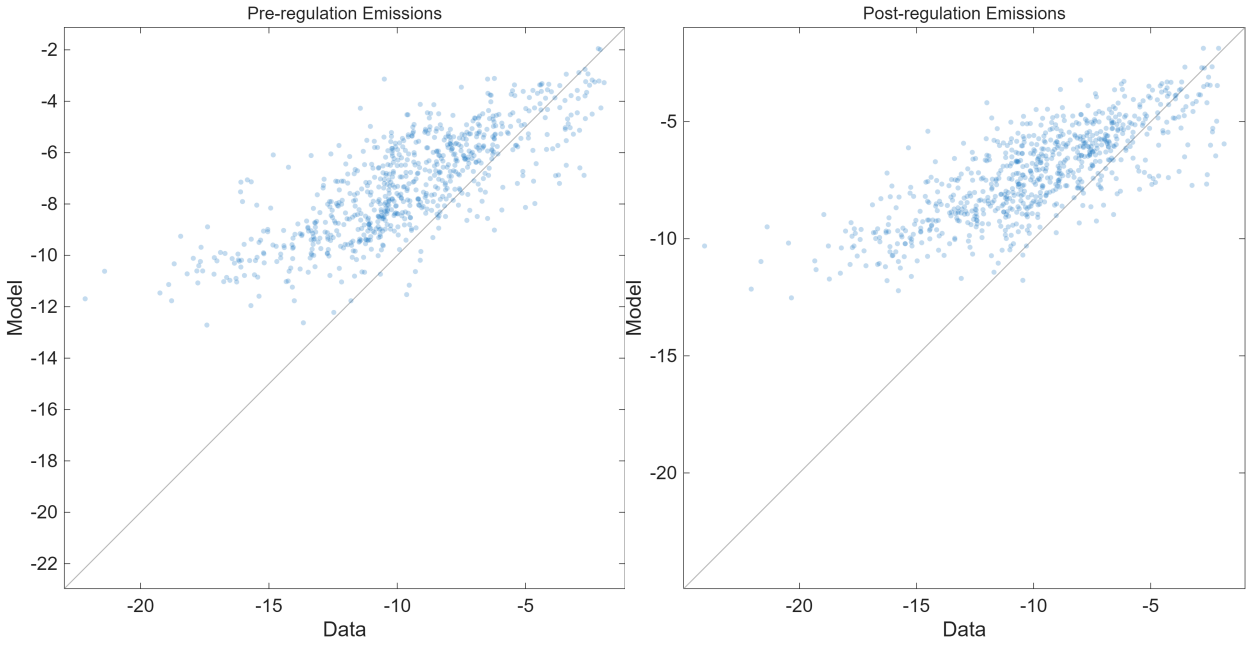
Notes: "All plants" includes plants belonging to both single-plant and multiplant firms; "Single-plant firms" restricts the sample to plants belonging to single-plant firms. The panels report the fuel-substitution elasticity λ and the firm-level demand elasticity σ, respectively. The outer Cobb–Douglas elasticity is fixed at one in both samples. Sectors 1–10 follow the order in Table 4. Bars show 95% confidence intervals.

Figure A.11: Model fit: production distributions



Notes: The panels compare pre-regulation manufacturing and power generation moments from the data and those of the model. They are targeted for calibration. Cell-level sales and profit moments are weighted by the corresponding firm or plant counts; the maximum-sales-share target consists of 10 sector-level means. Because the firm–plant incidence matrix cannot be released, its model counterpart is constructed from synthetic portfolios that preserve the released aggregate margins. The sales and profit panels plot the logarithms of normalized, count-weighted cell totals for cells with positive data and model values. The maximum-sales-share panel plots sector-level means in levels. The 45-degree line indicates a perfect fit between the model and the data. Profits are variable profits. There are 10 manufacturing sectors, 250 districts, and 4 power generation fuels.

Figure A.12: Model fit: emissions distributions



Notes: The panels compare pre- and post-regulation emissions distributions across pollutants and districts. These emissions distributions are not targeted in estimation and serve as validation moments. Each point represents a district–pollutant cell, combining manufacturing and power generation. For each pollutant and period, emissions are divided by their respective data or model totals over districts with positive values in both series, and the resulting shares are plotted in natural logs. The 45-degree line indicates a perfect fit between the model and the data. There are 3 pollutants and 250 districts.