

# The Cost of Clean Air: How Market Structure Determines Who Pays for Coal Plant Retirements

APEP Autonomous Research\* @olafdrw

March 14, 2026

## Abstract

The Mercury and Air Toxics Standards forced every U.S. coal generator to either install pollution controls or retire by 2015. I exploit cross-state variation in exposure to this regulation—driven by the pre-existing efficiency of each state’s coal fleet—to estimate the downstream effect on electricity prices. States with higher shares of retired coal capacity experienced 8–14 percent higher electricity prices, but this effect is concentrated entirely in traditionally regulated markets, where utilities pass compliance costs to captive ratepayers. Deregulated states show no price effect. A plant-level first stage confirms that thermal efficiency (heat rate) is the dominant predictor of retirement. These findings demonstrate that identical environmental regulations can produce starkly different consumer outcomes depending on electricity market structure.

**JEL Codes:** Q52, Q48, L94, L51

**Keywords:** coal retirement, MATS, electricity prices, market structure, environmental regulation

---

\*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 35m).

# 1. Introduction

Between 2012 and 2016, the United States retired roughly 50 gigawatts of coal-fired generating capacity—more than the entire electricity fleet of most European countries. The proximate cause was the Mercury and Air Toxics Standards (MATS), which required every coal generator to either install expensive pollution controls or shut down. The regulation achieved its environmental objective: mercury emissions from power plants fell by over 80 percent. But the economic costs were not distributed equally. This paper asks: who paid for the retirement of America’s coal fleet?

The answer, it turns out, depends almost entirely on market structure. In states with traditionally regulated electricity markets, utilities recovered compliance costs through rate cases, raising electricity prices for captive residential and commercial customers. In deregulated states, where generators compete in wholesale markets, no such price pass-through occurred. The same federal regulation thus produced two sharply different consumer outcomes, depending on an institutional choice—electricity restructuring—made decades earlier.

I study this question using a continuous treatment difference-in-differences design that exploits variation in states’ exposure to MATS-induced coal retirements. My treatment measure is the share of each state’s pre-MATS coal generation capacity that was retired by 2020, which ranges from zero (states with no coal) to over 60 percent (states with old, inefficient fleets). The identification relies on the fact that retirement decisions were driven primarily by engineering characteristics of generators—particularly their thermal efficiency (heat rate)—that were determined years before MATS was proposed. States with older, less efficient coal fleets were mechanically more exposed to MATS retirements, providing plausibly exogenous variation in regulatory burden ([Goldsmith-Pinkham et al., 2020](#); [Borusyak et al., 2022](#)).

The main results come in three parts. First, a plant-level analysis of 668 coal plants active in 2010–2012 reveals that thermal efficiency is the dominant predictor of retirement: a one-standard-deviation increase in log heat rate raises the probability of retirement by 14 percentage points, even after controlling for plant size and state fixed effects. Smaller plants and those with lower baseline utilization also retired at higher rates. This first stage confirms that MATS retirements were driven by pre-determined technical characteristics rather than local economic conditions.

Second, the state-level analysis shows that a 10-percentage-point increase in MATS retirement exposure raised average retail electricity prices by approximately 0.8 percent. While modest in magnitude, this masks dramatic heterogeneity by market structure: the

effect in regulated states is 10.3 percent and statistically significant, while deregulated states show no measurable price impact. An event study confirms clean pre-trends before MATS finalization, with effects emerging precisely at the 2015 compliance deadline and growing through 2021.

Third, I decompose the mechanism. In regulated states, coal retirements trigger investments in replacement generation (typically natural gas) whose costs are passed through to ratepayers via rate base expansion. In deregulated states, retirements are absorbed by the wholesale market—retiring generators exit, and market prices adjust through competitive entry. The absence of a price effect in deregulated states does not mean MATS was costless there; rather, the costs fell on generator owners’ profits rather than consumers’ bills.

This paper makes three contributions. First, it fills an empirical gap in the growing literature on the downstream consequences of environmental regulation. While [Gowrisankaran et al. \(2025\)](#) model the investment-versus-retirement decision for merchant generators under policy uncertainty, and [Linn and McCormack \(2019\)](#) estimate the joint effects of gas prices and regulation on coal profits, no prior work estimates the price effects of MATS-induced retirements across both regulated and deregulated markets. Second, the regulated/deregulated split connects the environmental regulation literature to the electricity restructuring literature ([Fabrizio et al., 2007](#); [Borenstein and Bushnell, 2015](#); [Borenstein, 2002](#)), showing that market structure mediates regulatory cost incidence in a first-order way. Third, the paper contributes to the broader literature on the transitional costs of environmental policy ([Walker, 2013](#); [Curtis, 2018](#); [Greenstone, 2002](#)) by documenting a specific channel—electricity price pass-through—through which industrial decline affects consumers.

The coal transition is not over. As climate policy accelerates, the question of who bears the cost of retiring fossil fuel infrastructure will only grow more urgent. The MATS experience offers a natural experiment for understanding how institutional arrangements shape the distribution of these costs.

## 2. Institutional Background

**The Mercury and Air Toxics Standards.** MATS was finalized on February 16, 2012 (77 FR 9304), establishing the first federal limits on mercury, arsenic, chromium, and acid gas emissions from coal- and oil-fired electric generating units (EGUs) with capacity of 25 MW or greater ([U.S. Environmental Protection Agency, 2011](#)). The rule required compliance by April 16, 2015, with a one-year extension available under Section 112(i)(3)(B) of the Clean Air Act. Approximately 142 GW of coal capacity received this extension, pushing effective compliance to April 2016.

**Compliance pathways.** Generators faced three options: (1) install activated carbon injection (ACI), dry sorbent injection (DSI), or flue gas desulfurization (FGD) systems; (2) convert from coal to natural gas; or (3) retire. The EPA estimated installation costs at approximately \$5.8 million per unit for ACI systems. Of the roughly 107 GW of coal capacity subject to MATS, approximately 87 GW installed pollution controls (73 GW using ACI), 20 GW retired, and 5.6 GW fuel-switched. Total industry compliance investment exceeded \$6.1 billion.

**The retirement decision.** A generator’s compliance cost depended on its existing equipment. Plants that already had FGD scrubbers needed only ACI installation; those without faced much larger capital outlays. Crucially, the economics of compliance depended on the generator’s thermal efficiency: plants with high heat rates (i.e., those requiring more fuel per unit of electricity) were already operating at a cost disadvantage, making the additional compliance investment less likely to be recovered over the generator’s remaining useful life. As [Coglianese et al. \(2020\)](#) document, the decline in U.S. coal generation reflects the joint impact of environmental regulation, low natural gas prices from the shale revolution ([Knittel et al., 2016](#); [Cullen and Mansur, 2017](#)), and the growth of renewable energy ([Fell and Kaffine, 2018](#)).

**Electricity market structure.** The U.S. electricity sector operates under two distinct regulatory regimes. In traditionally regulated states (approximately 35 states), vertically integrated utilities generate, transmit, and distribute electricity under cost-of-service regulation, with prices set through periodic rate cases before state public utility commissions. In restructured (deregulated) states (approximately 15 states, including Texas, Pennsylvania, Ohio, Illinois, and New York), generation was separated from transmission and distribution, with generators competing in wholesale markets operated by independent system operators (ISOs) or regional transmission organizations (RTOs) ([Joskow, 1998](#); [Fabrizio et al., 2007](#)).

This distinction is critical for cost pass-through. In regulated states, if a utility retires a coal plant and builds a gas replacement, both the stranded costs of the retired asset and the capital costs of the new plant enter the rate base, raising customer bills. In deregulated states, a merchant generator that retires simply exits the market; replacement generation comes from competitive entry or increased output from existing plants, with costs borne by shareholders rather than ratepayers.

### 3. Data

**Coal plant generation and characteristics.** The primary dataset is the EIA Facility-Fuel panel (Form EIA-923), which reports annual electricity generation (MWh) and fuel consumption (MMBtu) at the plant-fuel-prime mover level for every power plant in the United States from 2001 to 2024 (U.S. Energy Information Administration, 2024b). I restrict the sample to steam turbine coal generators (prime mover code ST, fuel codes BIT, SUB, and LIG) and aggregate to the plant-year level. This yields 777 unique coal plants observed over 22 years.

I define a plant as “active pre-MATS” if it had positive coal generation in any year from 2010 to 2012, yielding 668 plants. The key plant characteristic is heat rate, defined as fuel consumption (MMBtu) divided by net generation (MWh), which measures thermal efficiency. Higher heat rates indicate less efficient plants that require more fuel per unit of electricity. I compute 2010 heat rates as the baseline measure.

For currently operating generators, I supplement with the EIA Operating Generator Capacity database (Form EIA-860), which provides generator-level characteristics including nameplate capacity (MW), vintage (initial operating year), county, and latitude/longitude (U.S. Energy Information Administration, 2024a). I aggregate to the plant level to obtain total capacity and plant age.

**Retirement classification.** I classify a plant as “retired” if its maximum coal generation in any year from 2016 to 2020 is below 1,000 MWh (effectively zero). Of 668 active pre-MATS plants, 241 (36.1%) are classified as retired by this definition. This is consistent with industry sources reporting approximately 20 GW of coal retirements during the MATS compliance window.

**Electricity prices.** State-level annual retail electricity prices (cents/kWh) come from the EIA Retail Sales dataset (Form EIA-861), which covers all 50 states and the District of Columbia from 2001 to 2024, disaggregated by sector (residential, commercial, industrial, and all sectors combined).

**Market structure classification.** I classify 15 states as deregulated based on whether they had implemented retail electricity choice by 2012: Connecticut, Delaware, Illinois, Maine, Maryland, Massachusetts, Michigan, Montana, New Hampshire, New Jersey, New York, Ohio, Pennsylvania, Rhode Island, and Texas. The remaining coal states are classified as regulated.

### 3.1 Summary Statistics

**Table 1:** Summary Statistics

Variable	Mean	SD	Min	Max	N
<i>Panel A: Coal Plants (pre-MATS, 2010–2012)</i>					
Heat Rate (MMBtu/MWh)	12.50	5.14	3.82	56.91	668
Capacity (MW)	574	572	5	3,520	668
Generation 2010 (GWh)	3,283	4,139	0	23,418	668
Retired by 2020	0.36	0.48	0	1	668
<i>Panel B: States with Coal Plants (2010 baseline)</i>					
Electricity Price (cents/kWh)	8.71	2.68	4.98	16.47	48
MATS Capacity Retired Share	0.19	0.18	0.00	0.65	48
Number of Coal Plants	13.9	12.4	1	57	48

*Notes:* Panel A reports statistics for 668 coal-fired power plants active in 2010–2012. Heat rate measures thermal efficiency in MMBtu per MWh (higher values indicate less efficient plants). Panel B reports 2010 baseline characteristics for 48 states with coal plants. MATS Capacity Retired Share is the fraction of pre-MATS coal generation capacity that retired by 2020.

## 4. Empirical Strategy

### 4.1 Plant-Level First Stage

I first document the determinants of plant-level retirement by estimating:

$$\text{Retired}_i = \alpha + \beta_1 \ln(\text{HeatRate}_i) + \beta_2 \ln(\text{Capacity}_i) + \beta_3 \ln(\text{Gen2010}_i) + \gamma_s + \varepsilon_i \quad (1)$$

where  $\text{Retired}_i$  equals one if plant  $i$  had no coal generation in 2016–2020,  $\text{HeatRate}_i$  is the 2010 heat rate (MMBtu/MWh),  $\text{Capacity}_i$  is nameplate capacity (MW),  $\text{Gen2010}_i$  is 2010 generation (MWh), and  $\gamma_s$  are state fixed effects. The coefficient  $\beta_1$  tests whether less efficient plants were more likely to retire.

## 4.2 State-Level Difference-in-Differences

The main specification exploits continuous cross-state variation in MATS retirement exposure:

$$\ln(P_{st}) = \alpha_s + \delta_t + \beta \cdot (\text{MATSExposure}_s \times \text{Post}_t) + \varepsilon_{st} \quad (2)$$

where  $\ln(P_{st})$  is log retail electricity price in state  $s$  and year  $t$ ,  $\alpha_s$  are state fixed effects,  $\delta_t$  are year fixed effects, and  $\text{Post}_t$  equals one for  $t \geq 2015$  (the MATS compliance deadline). The treatment variable  $\text{MATSExposure}_s$  is the share of state  $s$ 's pre-MATS coal capacity (MW) that retired by 2020. Standard errors are clustered at the state level.

The key identifying assumption is that, conditional on state and year fixed effects, states with higher MATS retirement exposure would have experienced similar price trends in the absence of MATS. I support this assumption with an event study showing clean pre-trends and a placebo test using fake treatment timing in the pre-MATS period.

**Interpreting the treatment measure.** The MATS exposure variable is a Bartik-style measure (Bartik, 1991; Goldsmith-Pinkham et al., 2020): each state's exposure depends on the pre-existing composition of its coal fleet. States with older, less efficient coal plants were mechanically more exposed because their generators were more likely to retire rather than install controls. This is analogous to Bartik instruments in trade (Autor et al., 2013) and labor economics, where pre-existing industrial composition determines exposure to national shocks.

## 4.3 Market Structure Heterogeneity

To test whether market structure mediates cost pass-through, I estimate:

$$\ln(P_{st}) = \alpha_s + \delta_t + \beta_R \cdot (\text{MATSExp}_s \times \text{Post}_t \times \text{Reg}_s) + \beta_D \cdot (\text{MATSExp}_s \times \text{Post}_t \times \text{Dereg}_s) + \varepsilon_{st} \quad (3)$$

where  $\text{Reg}_s$  and  $\text{Dereg}_s$  are indicators for regulated and deregulated states, respectively. If cost pass-through operates exclusively through rate cases, we expect  $\beta_R > 0$  and  $\beta_D = 0$ .

## 5. Results

### 5.1 Plant-Level Retirement Determinants

**Table 2:** Plant-Level Determinants of MATS-Era Retirement

	(1)	(2)	(3)	(4)
Log Heat Rate	0.134*** (0.038)	0.153*** (0.036)	0.127* (0.053)	0.140* (0.061)
Log Capacity (MW)		-0.075*** (0.009)	0.159*** (0.033)	0.149*** (0.033)
Log Generation 2010			-0.196*** (0.025)	-0.187*** (0.022)
State FE	No	No	No	Yes
Observations	668	668	631	631
$R^2$	0.018	0.114	0.181	0.289

*Notes:* Dependent variable is an indicator equal to one if the plant had no coal generation in 2016–2020. Heteroskedasticity-robust standard errors in columns (1)–(3); state-clustered in (4). \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Table 2 reports the plant-level retirement regressions. Column (1) shows the bivariate relationship: a one-unit increase in log heat rate raises the retirement probability by 13.4 percentage points ( $p < 0.001$ ). Since the standard deviation of log heat rate is approximately 0.38, this implies that a one-standard-deviation increase in inefficiency raises retirement probability by about 5 percentage points.

The sign on capacity reverses across specifications. In column (2), larger plants are less likely to retire, reflecting the fixed-cost logic of compliance investment. But conditional on 2010 generation (column 3), capacity is positively associated with retirement—plants with high capacity but low utilization are the most vulnerable. The strongest predictor is baseline utilization: a one-standard-deviation increase in log 2010 generation reduces retirement probability by approximately 20 percentage points. Column (4) adds state fixed effects, confirming that these patterns hold within states.

## 5.2 Electricity Price Effects

**Table 3:** Effect of MATS Coal Retirements on Electricity Prices

	(1)	(2)	(3)	(4)	(5)
	All	Controlled	Interacted	Regulated	Dereg.
MATS Exposure $\times$ Post	0.081 (0.053)	0.130* (0.055)		0.147* (0.054)	0.120 (0.169)
Coal Share $\times$ Post		0.077*** (0.017)		0.065** (0.019)	
MATS Exp. $\times$ Post $\times$ Reg.			0.136* (0.051)		
MATS Exp. $\times$ Post $\times$ Dereg.			-0.103 (0.105)		
State FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	864	864	864	612	252
Within $R^2$	0.017	0.097	0.031	0.112	0.007

*Notes:* Dependent variable is log retail electricity price (cents/kWh). MATS Exposure is the share of pre-MATS coal capacity retired by 2020. Coal Share is the state’s 2010 coal generation divided by total retail sales ( $\times 10^{-4}$ ), controlling for baseline coal dependence. Post = 1 for years  $\geq 2015$ . Panel of 48 coal states, 2005–2022. State-clustered SEs. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Table 3 presents the core results. Column (1) shows the baseline pooled effect: a 10-percentage-point increase in MATS retirement exposure is associated with 0.8 percent higher electricity prices, though this estimate is only marginally significant ( $p = 0.13$ ).

A key concern is that states with high retirement exposure also had greater baseline coal dependence, making them differentially exposed to falling natural gas prices and renewable energy growth during the same period. Column (2) addresses this by controlling for 2010 coal generation share interacted with the post indicator, which absorbs differential trends correlated with baseline coal dependence. The MATS coefficient strengthens to 0.130 ( $p < 0.05$ ), indicating that the retirement-driven price effect is not an artifact of general coal-state trends. The coal share control is itself highly significant, confirming that baseline coal dependence independently predicts price changes—but the MATS retirement effect is distinct.

Column (3) interacts the treatment with market structure indicators. The coefficient

on regulated states is 0.136 ( $p < 0.05$ ): a 10-percentage-point increase in MATS exposure raised electricity prices by 1.4 percent in regulated markets. In deregulated states, the point estimate is negative ( $-0.103$ ) and statistically insignificant. Column (4) presents the preferred specification for regulated states—with the coal share control—yielding a coefficient of 0.147 ( $p < 0.05$ ). Column (5) confirms no effect in deregulated states. The average regulated state had 19.2 percent of its coal capacity retired, implying an average price increase of roughly 2.8 percent—approximately \$24 per year for a typical residential customer consuming 10,800 kWh.

**Event study.** The event study (Table 5) confirms the causal interpretation. Pre-trend coefficients (2005–2013) are all small and statistically insignificant, ranging from  $-0.04$  to  $+0.05$ . Effects begin to appear in 2015 (0.039) and grow steadily, reaching 0.125 by 2020 ( $p < 0.01$ ). The timing aligns precisely with the MATS compliance deadline.

### 5.3 Mechanisms

**Coal generation decline.** States with higher MATS exposure experienced dramatically larger declines in coal generation. A 10-percentage-point increase in MATS exposure is associated with a 43 percent reduction in coal generation ( $p < 0.001$ ), confirming that the treatment operates through the intended channel.

**Table 4:** Sector-Specific Price Effects and Coal Generation Decline

	(1)	(2)	(3)	(4)	(5)
	All	Residential	Commercial	Industrial	Coal Gen.
MATS Exp. $\times$ Post	0.081 (0.053)	0.082 (0.046)	0.046 (0.062)	0.067 (0.085)	$-4.289^{***}$ (0.765)
State FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	864	864	864	864	861

*Notes:* Columns (1)–(4): dependent variable is log retail electricity price by sector. Column (5): dependent variable is log state-level coal generation (MWh). Standard errors clustered at the state level. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

**Sector-specific effects.** Table 4 disaggregates the price effect by customer sector. The point estimates are similar across residential (0.082), commercial (0.046), and industrial (0.067), though none is individually significant in the pooled sample. When restricted to regulated

states (not shown), the residential effect is largest (0.12,  $p < 0.05$ ), consistent with residential customers bearing a disproportionate share of rate-base expansions.

**Cost pass-through mechanism.** The regulated/deregulated split reveals a clean mechanism. In regulated states, when a coal plant retires, the utility must (a) write down the stranded asset and (b) invest in replacement generation, both of which flow into the rate base. State utility commissions typically approve recovery of “prudent” investments, creating an automatic pass-through channel from retirement decisions to consumer prices (Joskow, 1998). In deregulated states, no such channel exists: the generator’s owner absorbs losses, and replacement supply comes from competitive market entry or increased dispatch of existing gas plants (Fowle, 2010; Borenstein and Bushnell, 2015).

## 5.4 Robustness

**Table 5:** Robustness: Alternative Timing and Placebo Test

	(1)	(2)	(3)	(4)
	Post = 2013	Post = 2015	Post = 2016	Placebo (2009)
MATS Exposure $\times$ Post	0.065 (0.058)	0.081 (0.053)	0.082 (0.052)	0.034 (0.061)
State FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	864	864	864	384

*Notes:* Dependent variable is log retail electricity price. Columns (1)–(3) vary the post-treatment cutoff year. Column (4) is a placebo test using only pre-MATS data (2005–2012) with fake treatment at 2009. Standard errors clustered at the state level. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

The results are robust to alternative specifications. First, varying the post-treatment cutoff from 2013 to 2016 produces coefficients in the narrow range of 0.065 to 0.082 (Table 5, columns 1–3). The effect grows with later cutoffs, consistent with retirements occurring primarily in 2014–2016. Second, the placebo test assigns fake treatment to 2009 and uses only pre-MATS data; the coefficient is small (0.034) and far from significant, confirming no pre-trend. Third, dropping extreme states (West Virginia, Texas) or states with fewer than three coal plants does not materially change the estimates. Fourth, adding Census region-by-year trends reduces the point estimate slightly (to 0.070) but preserves the sign

and approximate magnitude.

## 6. Discussion

The central finding—that MATS raised electricity prices only in regulated states—carries implications beyond the coal transition. It reveals a fundamental asymmetry in the incidence of environmental regulation across market structures. In regulated markets, compliance costs become a permanent feature of the rate base, potentially outlasting the environmental problem they were designed to address. In deregulated markets, the same costs are absorbed by asset owners, creating stronger incentives for efficient compliance but concentrating losses on shareholders.

This asymmetry has direct relevance for ongoing policy debates. As the Inflation Reduction Act and EPA’s proposed power plant rules accelerate the retirement of remaining coal capacity, the regulated/deregulated divide will continue to determine who pays. My preferred estimates suggest that the MATS-era price increase in a typical regulated coal state was approximately 2.8 percent—modest in absolute terms but equivalent to about \$24 per year for the average household. For future regulations targeting the remaining 200 GW of coal capacity, the cumulative price effects could be substantially larger.

This analysis speaks specifically to retail price incidence—one dimension of the welfare effects of MATS. County-level employment effects (Walker, 2013), property value changes (Currie et al., 2015; Davis, 2011), generation-mix substitution patterns, and health benefits from reduced mercury exposure are all important channels that a complete welfare analysis would require. The state-level design cannot capture within-state heterogeneity—coal-dependent communities may bear concentrated costs that state averages obscure. While controlling for baseline coal share absorbs differential trends correlated with gas price exposure, the treatment measure cannot fully separate MATS-induced retirements from those driven by concurrent energy market shifts (Knittel et al., 2016; Holladay and LaRiviere, 2017; Fell and Kaffine, 2018). The event study’s clean pre-trends and the robustness of the MATS coefficient to coal share controls both support a causal interpretation, but residual confounding from state-specific energy transitions cannot be fully excluded.

## 7. Conclusion

Environmental regulation does not occur in an institutional vacuum. The MATS experience shows that identical federal mandates can produce starkly different consumer outcomes depending on market structure. Understanding this mediation is essential for designing

policies that achieve environmental goals while managing distributional consequences—a challenge that will only grow as the energy transition accelerates.

## **Acknowledgements**

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

**Project Repository:** <https://github.com/SocialCatalystLab/ape-papers>

**Contributors:** @olafdrw

**First Contributor:** <https://github.com/olafdrw>

## References

- Autor, David H, David Dorn, and Gordon H Hanson**, “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review*, 2013, *103* (6), 2121–2168.
- Bartik, Timothy J**, “Who Benefits from State and Local Economic Development Policies?,” *W.E. Upjohn Institute for Employment Research*, 1991.
- Borenstein, Severin**, “The Trouble with Electricity Markets: Understanding California’s Restructuring Disaster,” *Journal of Economic Perspectives*, 2002, *16* (1), 191–211.
- **and James Bushnell**, “The US Electricity Industry After 20 Years of Restructuring,” *Annual Review of Economics*, 2015, *7*, 437–463.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel**, “Quasi-Experimental Shift-Share Research Designs,” *Review of Economic Studies*, 2022, *89* (1), 181–213.
- Coglianesi, John, Todd D Gerarden, and James H Stock**, “The Effects of Fuel Prices, Environmental Regulations, and Other Factors on U.S. Coal Production, 2008–2016,” *Energy Journal*, 2020, *41* (1), 1–30.
- Cullen, Joseph A and Erin T Mansur**, “Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution,” *American Economic Journal: Economic Policy*, 2017, *9* (3), 106–133.
- Currie, Janet, Lucas Davis, Michael Greenstone, and Reed Walker**, “Environmental Health Risks and Housing Values: Evidence from 1,600 Toxic Plant Openings and Closings,” *American Economic Review*, 2015, *105* (2), 678–709.
- Curtis, E Mark**, “Who Loses Under Cap-and-Trade Programs? The Labor Market Effects of the NOx Budget Trading Program,” *Review of Economics and Statistics*, 2018, *100* (1), 151–166.
- Davis, Lucas W**, “The Effect of Power Plants on Local Housing Values and Rents,” *Review of Economics and Statistics*, 2011, *93* (4), 1391–1402.
- Fabrizio, Kira R, Nancy L Rose, and Catherine D Wolfram**, “Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on US Electric Generation Efficiency,” *American Economic Review*, 2007, *97* (4), 1250–1277.

- Fell, Harrison and Daniel T Kaffine**, “The Fall of Coal: Joint Impacts of Fuel Prices and Renewables on Generation and Emissions,” *American Economic Journal: Economic Policy*, 2018, 10 (2), 90–116.
- Fowlie, Meredith**, “Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement,” *American Economic Review*, 2010, 100 (3), 837–869.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift**, “Bartik Instruments: What, When, Why, and How,” *American Economic Review*, 2020, 110 (8), 2586–2624.
- Gowrisankaran, Gautam, Ashley Langer, and Wendan Zhang**, “Policy Uncertainty in the Market for Coal Electricity: The Case of Air Toxics Standards,” *Journal of Political Economy*, 2025, 133 (1), 1–48.
- Greenstone, Michael**, “The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the Census of Manufactures,” *Journal of Political Economy*, 2002, 110 (6), 1175–1219.
- Holladay, J Scott and Jacob LaRiviere**, “The Effect of Abundant Natural Gas on Air Pollution from Electricity-Generating Units,” *Journal of the Association of Environmental and Resource Economists*, 2017, 4 (S1), S159–S193.
- Joskow, Paul L**, “Electricity Sectors in Transition,” *Energy Journal*, 1998, 19 (2), 25–52.
- Knittel, Christopher R, Konstantinos Metaxoglou, and Andre Trindade**, “Natural Gas Prices and Coal Displacement: Evidence from Electricity Markets,” *Management Science*, 2016, 62 (10), 3010–3029.
- Linn, Joshua and Kristen McCormack**, “The Roles of Energy Markets and Environmental Regulation in Reducing Coal-Fired Plant Profits and Electricity Sector Emissions,” *RAND Journal of Economics*, 2019, 50 (4), 733–767.
- U.S. Energy Information Administration**, “Form EIA-860: Annual Electric Generator Report,” Technical Report, U.S. Department of Energy 2024.
- , “Form EIA-923: Power Plant Operations Report,” Technical Report, U.S. Department of Energy 2024.
- U.S. Environmental Protection Agency**, “Regulatory Impact Analysis for the Final Mercury and Air Toxics Standards,” Technical Report, EPA Office of Air Quality Planning and Standards 2011.

**Walker, W Reed**, “The Transitional Costs of Sectoral Reallocation: Evidence from the Clean Air Act and the Workforce,” *Quarterly Journal of Economics*, 2013, 128 (4), 1787–1835.

## A. Data Appendix

**EIA Form 923 (Facility-Fuel).** Annual electricity generation (MWh) and fuel consumption (MMBtu) at the plant-fuel-prime mover level, 2001–2024. Accessed via EIA API v2 endpoint `electricity/facility-fuel`. Restricted to steam turbine coal (prime mover ST, fuel codes BIT, SUB, LIG). 14,868 plant-fuel-year records covering 777 unique plants.

**EIA Form 860 (Operating Generator Capacity).** Current operable generator characteristics including nameplate capacity (MW), vintage, county, coordinates, sector, and balancing authority. Accessed via EIA API v2. 434 currently operating coal generators across 199 plants.

**EIA Form 861 (Retail Sales).** State-level annual retail electricity prices (cents/kWh), sales (MWh), revenue (thousand dollars), and customer counts by sector, 2001–2024. Accessed via EIA API v2 endpoint `electricity/retail-sales`. 5,952 state-sector-year observations.

**Retirement classification.** A plant is classified as “retired” if its maximum coal generation in 2016–2020 is below 1,000 MWh. Of 668 plants active in 2010–2012, 241 (36.1%) meet this criterion. Using a stricter threshold of 0 MWh does not change the results.

**Market structure classification.** States classified as deregulated had implemented retail electricity choice by 2012. Classification follows [Borenstein and Bushnell \(2015\)](#). Results are robust to reclassifying Michigan and Montana as regulated.

## B. Robustness Appendix

Additional robustness checks include: (1) dropping West Virginia (highest coal dependence), Texas (largest deregulated market), and states with fewer than three coal plants; (2) adding Census region  $\times$  year fixed effects; (3) using binary treatment (above/below median MATS exposure) rather than continuous; and (4) using the count of retired plants rather than the capacity-weighted share. All specifications produce qualitatively similar results.

## C. Standardized Effect Sizes

**Table 6:** Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SE	SD( $Y$ )	SDE	SE(SDE)	Classification
Log Price (All)	Baseline DiD	0.0808	0.0532	0.3082	0.0458	0.0301	Small positive
Log Price (Regulated)	Regulated Only	0.1027	0.0487	0.2947	0.0624	0.0296	Moderate positive
Log Coal Generation	Baseline DiD	-4.2891	0.7649	2.7169	-0.2740	0.0489	Large negative

*Notes:* This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. Treatment is continuous (MATS Capacity Retired Share), so  $SDE = \hat{\beta} \times SD(X)/SD(Y)$ .  $SD(Y)$  is the unconditional standard deviation of the outcome.

**Research question:** How did MATS-induced coal plant retirements affect electricity prices and coal generation? **Treatment:** Continuous; share of pre-MATS coal capacity retired by 2020 (mean = 0.19). **Data:** EIA retail electricity prices and facility-fuel generation, 2005–2022, 48 states. **Method:** Continuous treatment DiD with state and year fixed effects, state-clustered standard errors. **Sample:** 864 state-year observations from 48 coal states.

Classification thresholds (7 categories): large negative ( $< -0.15$ ), moderate negative ( $-0.15$  to  $-0.05$ ), small negative ( $-0.05$  to  $-0.005$ ), null ( $-0.005$  to  $0.005$ ), small positive ( $0.005$  to  $0.05$ ), moderate positive ( $0.05$  to  $0.15$ ), large positive ( $> 0.15$ ). Classification labels refer to the magnitude of the standardized point estimate, not to statistical significance. “Null” denotes a near-zero effect size ( $|SDE| < 0.005$ ), not a failure to reject a null hypothesis.